

Windows® IT Pro

MAY 2009 | WINDOWSITPRO.COM | WE'RE IN IT WITH YOU

DNS

What's New in Name Resolution? p. 15

Windows
Server 2008 R2 p. 23

Organize Active Directory
Objects p. 27

Track Printer Changes p. 35


Create PowerShell Functions p. 39

Move Public Folders
to SharePoint p. 45

Buyer's Guide:
KVM over IP
Switches p. 55



Microsoft's
Doug Leland




THE BARRIERS TO VIRTUALIZATION FALL AWAY.

THE VIRTUAL. THE PHYSICAL.

YOU COMMAND THEM BOTH.

HOW FAR WILL YOU TAKE VIRTUAL?

With Microsoft® Virtualization you can manage both physical and virtual Hyper-V™ servers and desktops using the same management platform, Microsoft System Center. And if you have VMware ESX, System Center can manage it, too. From the datacenter to the desktop, you manage virtualization more powerfully and easily than ever before. Explore the virtual at microsoft.com/virtualization



To get the full story on your phone, snap a picture of this tag. (Requires a free mobile app from <http://gettag.mobi>)

Microsoft | Virtualization



COVER STORY

15 What's New in DNS and Name Resolution?

Windows Server 2008 and Windows Vista have introduced some important changes to Windows name resolution and DNS. If you're hoping to finally achieve a WINS-free environment, you need to know about these changes.

BY MARK MINASI

FEATURES

19 Identity and Security: Microsoft's Next Generation

Microsoft's new Identity and Security Business Group could make it easier for IT pros to deploy and manage their access and security infrastructures.

BY JEFF JAMES

23 Inside Windows Server 2008 R2

Big changes in Server 2008 R2 include Live Migration, the Active Directory Administrative Center (ADAC), and the PowerShell Integrated Scripting Environment (ISE).

BY MICHAEL OTEY

27 Organize Your Active Directory Objects

You can control Active Directory objects better when they're tied to a guardian you set up.

BY TONY MURRAY

31 A First Look at Exchange 2010

Exchange Server 2010 includes significant changes in the Information Store, a new approach to high availability through database replication, management and administration updates, and new features for messaging compliance.

BY TONY REDMOND

35 Scripting Utilities to Keep Tabs on Your Printers

Tracking changes to printer configurations in a large organization can be very difficult, but these two scripts can save you troubleshooting effort by recording changes and letting you compare data from different dates.

BY JIM TURNER

39 Create Your Own PowerShell Functions

You can make PowerShell functions as simple or as complex as necessary. Here's what you need to know to get started.

BY ROBERT SHELDON

OFFICE & SHAREPOINT PRO

45 Moving your Public Folders to SharePoint

Migrating your Exchange public folders to SharePoint is a complex task, but with this guidance you can avoid many of the pitfalls.

BY RON CHARITY

INTERACT

12 Reader to Reader

Encrypt a single file or a bunch of files with AxCrypt, and use DevCon to manage devices from the command line.

13 Ask the Experts

Find out how to join a computer to a domain in Windows 7, where to put your domain controllers in a virtualized environment, and how to see the full header of your emails in Outlook.

PRODUCTS

51 New & Improved

Check out the latest products to hit the marketplace.

PRODUCT SPOTLIGHT: Citrix Essentials for Hyper-V and XenServer

REVIEW

53 Xobni

This easy-to-use Outlook add-on provides enhanced search functionality.

BY ANNE GRUBB

REVIEW

53 PatchSee

Simple but useful Ethernet cables can help you see the light when it comes to figuring out which end leads where.

BY ZAC WIGGY

REVIEW

54 Unbounded Printing Services for SharePoint

Tired of SharePoint's printing limitations? This product offers enhanced print services that make it easier to print SharePoint documents.

BY CURT SPANBURGH

BUYER'S GUIDE

55 KVM over IP Switches

Using a KVM over IP switch, one person can easily manage systems in several locations.

BY JASON BOVBERG

57 Industry Bytes

Intel is pushing back its Itanium processor while AMD pushes into the enterprise; what to do if you're blindsided by a layoff; the top nine skills for 2009; and what one study says about the importance of unified communications, and why it might not be accurate.

Windows IT Pro

A PENTON PUBLICATION

MAY_2009
VOLUME_15
NO_5

COLUMNS

JAMES | IT PRO PERSPECTIVE



3 Are You Turning to Virtualization to Cut IT Expenses?

Virtualization can help save money and increase efficiency in cash-strapped IT environments.

THURROTT | NEED TO KNOW



6 What You Need to Know About Windows 7 RC

The Windows 7 release candidate gives you a near-final look at what might be Microsoft's best OS to date.

MINASI | WINDOWS POWER TOOLS



8 2 Useful Bcdedit Options

Using Bcdedit, you can enable or disable DEP in certain situations, and you can determine how many CPUs Windows uses.

OTey | TOP 10



9 Windows Vista Shortcomings as a Business OS

Well-publicized problems such as speed and compatibility combine with Wi-Fi network connection problems and other glitches to prevent Windows Vista from providing a business ROI.

MORALES | WHAT WOULD MICROSOFT SUPPORT DO?



10 Conquer Desktop Heap Problems

Getting too many out of memory messages or application startup failures? Such problems are symptoms of desktop heap issues, which you can troubleshoot and solve using Task Manager and the Dheapmon tool and registry settings.

Access articles online at www.windowsitpro.com. Enter the article ID (located at the end of each article) in the InstantDoc ID text box on the home page.

IN EVERY ISSUE



- 4** letters@
windowsitpro.com
- 5** Your Savvy Assistant
- 63** Directory of Services
- 63** Advertising Index
- 63** Vendor Directory
- 64** Ctrl+Alt+Del



ON THE WEB

Read these articles at www.windowsitpro.com.

Top 5 Best Operations Manager Extensions

These five extensions let you do more with Systems Center Operations Manager 2007, including view the health of your servers and applications in diagram format. —Cameron Fuller

InstantDoc ID 101719

Five Ways to Manage Server Core

Managing Server Core requires you to dust off your command-line skills. Learn one local and four remote methods of management. —J. Peter Bruzzese

InstantDoc ID 101710

Script Trek

The Star Trek crew discovers how bad assumptions about data types can cause errors in VBScript and PowerShell code. —Dimitrios Kalemis

InstantDoc ID 101722

Corrupted boot.ini file Might be the Culprit When You Can't Find an Existing Windows Installation

A corrupted boot.ini file can bring your effort to fix an existing Windows installation to a screeching stop. Here's how to repair a boot.ini file so that it can correctly locate existing installations. —Oguzhan Oguz

InstantDoc ID 101723

New Ways to Reach Windows IT Pro Editors

Twitter: Visit the *Windows IT Pro* Twitter page at www.twitter.com/windowsitpro.

LinkedIn: To check out the *Windows IT Pro* group on LinkedIn, sign in on the LinkedIn homepage (www.linkedin.com), select the Search Groups option from the pull-down menu, and use "Windows IT Pro" as your search term.

Facebook: We've created a page on Facebook for *Windows IT Pro*, which you can access at <http://tinyurl.com/d5bquf>. Visit our Facebook page to read the latest reader comments, see links to our latest web content, browse our classic cover gallery, and participate in our Facebook discussion board.

Windows IT Pro

EDITORIAL

Editorial and Custom Strategy Director
Michele Crockett mcrockett@windowsitpro.com

Editor-in-Chief, Web Content Strategist
Jeff James jjames@windowsitpro.com

Executive Editor, IT Group
Amy Eisenberg amy@windowsitpro.com

Technical Director
Michael Otey motey@windowsitpro.com

Custom Group Editorial Director
Dave Bernard dbernard@windowsitpro.com

Web and Developer Strategic Editor
Anne Grubb agrubb@windowsitpro.com

Systems Management
Karen Bemowski kbemowski@windowsitpro.com
Caroline Marwitz cmarwitz@windowsitpro.com
Zac Wiggy zwiggy@windowsitpro.com

Messaging, Mobility, SharePoint, and Office
Brian Keith Winstead bwinstead@windowsitpro.com

Networking and Hardware
Jason Bovberg jbovberg@windowsitpro.com

Security
Lavon Peters lpeters@windowsitpro.com

SQL Server
Megan Bearly Keller mkeller@windowsitpro.com
Sheila Molnar smolnar@windowsitpro.com

Production Editor
Brian Reinholz breinholz@windowsitpro.com

IT Media Group Editors
Linda Harty, Chris Maxcer, David Riggs, Rita-Lyn Sanders

CONTRIBUTORS

News Editor
Paul Thurrott news@windowsitpro.com

SharePoint and Office Community Editor
Dan Holme danh@intelliem.com

Senior Contributing Editors
David Chernicoff david@windowsitpro.com
Mark Joseph Edwards mje@windowsitpro.com
Kathy Ivens kivens@windowsitpro.com
Mark Minasi mark@minasi.net
Paul Robichaux paul@robichaux.net
Mark Russinovich mark@sysinternals.com

Contributing Editors
Alex K. Angelopoulos aka@mvps.org
Sean Deuby sdeuby@windowsitpro.com
Michael Dragone mike@mikerochip.com
Jeff Felling jeff@blackstatic.com
Brett Hill brett@iisanswers.com
Darren Mar-Elia dmarella@windowsitpro.com
Tony Redmond tony.redmond@hp.com
Ed Roth eroth@windowsitpro.com
Eric B. Rux ericbrux@whshelp.com
William Sheldon bsheldon@interknowlogy.com
Randy Franklin Smith rsmith@montereytechgroup.com
Curt Spanburgh cspanburgh@scg.net
Orin Thomas orin@windowsitpro.com
Douglas Toombs help@toombs.us
Ethan Wilansky ewilansky@windowsitpro.com

ART & PRODUCTION

Senior Art Director
Larry Purvis lpurvis@windowsitpro.com

Art Director
Layne Petersen layne@windowsitpro.com

Production Director
Linda Kirgesler linda@windowsitpro.com

Senior Production Manager
Kate Brown kbrown@windowsitpro.com

Assistant Production Manager
Erik Lodermeier erik.lodermeier@penton.com

ADVERTISING SALES

Publisher Peg Miller
pmiller@windowsitpro.com

EMEA Managing Director Irene Clapham
irene.clapham@penton.com

Director of Sales Birdie J. Ghiglione
birdie.ghiglione@penton.com, 619-442-4064

Online Sales and Marketing Manager Dina Baird
Dina.Baird@penton.com

Key Account Directors Jeff Carnes jeff.carnes@penton.com
678-455-6146

Chrissy Ferraro christina.ferraro@penton.com
970-203-2883

Jacquelyn Baillie jacquelyn.baillie@penton.com
714-623-5007

Account Executives Barbara Ritter barbara.ritter@penton.com
858-759-3377
Cass Schulz cassandra.schulz@penton.com
858-357-7649

Client Project Managers Michelle Andrews 970-613-4964
Kim Eck 970-203-2953

Ad Production Supervisor Glenda Vaught glenda.vaught@penton.com

MARKETING & CIRCULATION

Customer Service 800-793-5697 (US and Canada)
44-161-929-2800 (Europe)

IT Group Audience Development Director Marie Evans marie.evans@penton.com

Marketing Director Sandy Lang sandy.lang@penton.com

CORPORATE



Penton Media, Inc.

Chief Executive Officer Sharon Rowlands Sharon.Rowlands@penton.com

Chief Financial Officer/Executive Vice President Jean Clifton jean.clifton@penton.com

TECHNOLOGY GROUP

Senior Vice President, Technology Media Group Kim Paulsen kpaulsen@windowsitpro.com

Windows®, Windows Vista®, and Windows Server® are trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries and are used by Penton Media under license from owner. *Windows IT Pro* is an independent publication not affiliated with Microsoft Corporation.

WRITING FOR WINDOWS IT PRO

Submit queries about topics of importance to Windows managers and systems administrators to articles@windowsitpro.com.

PROGRAM CODE

Unless otherwise noted, all programming code in this issue is © 2009, Penton Media, Inc., all rights reserved. These programs may not be reproduced or distributed in any form without permission in writing from the publisher. It is the reader's responsibility to ensure procedures and techniques used from this publication are accurate and appropriate for the user's installation. No warranty is implied or expressed.

LIST RENTALS

Contact Walter Karl, Inc. at 2 Blue Hill Plaza, 3rd Floor, Pearl River, NY 10965 or www.walterkarl.com/mailings/pentonLD/index.html.

REPRINTS

Diane Madzelonka, Diane.madzelonka@penton.com, 216-931-9268, 888-858-8851

"Recent economic pressures are driving growth of newer types of virtualization."



Are You Turning to Virtualization to Cut IT Expenses?

Minimize infrastructure costs and improve efficiency

I used to enjoy listening to the radio during my daily drive to work. Not so much anymore: Every newscast seems to offer a gloomier economic outlook than the last one, filled with news of layoffs, declining consumer confidence, and a host of other financial ills. In addition, everyone I know is cutting back on expenses, from my next-door neighbor to the corporate executives here at Penton Media. IT professionals haven't escaped the economic downturn, but some time-worn IT technologies are increasingly being called upon to minimize infrastructure costs and improve efficiency. Virtualization is at the forefront of this cost savings, and IT pros are turning to it more than ever before.

I recently spoke with Bob Meyer, Worldwide Virtualization Solutions Lead for the Technology Solutions Group at HP, about the effect virtualization is having on IT departments. Meyer told me that the economic climate is already having an effect (albeit a positive one) on HP's business, as customers who might have been dragging their feet on adopting a virtualization solution have moved ahead. "We've seen significant changes in customer behavior in the last 6 months," he said. "The economic climate has helped push some of our customers—who may have been on the fence when it comes to using virtualization for production environments—into using virtualization more aggressively."

Although server consolidation and testing environments have always been prime candidates for virtualization, recent economic pressures are driving growth of newer types of virtualization, including virtual desktop infrastructure (VDI) and virtualized desktops. "Virtualization [at the desktop] is one of the fastest growing areas of virtualization," says Meyer. "Virtualizing everything from the desktop to the data center provides lots of cost and efficiency savings."

Thin Clients, Fat Rewards

Although thin-client computing has been around for ages, recent virtualization developments promise to make this model more cost effective and usable than in the past. For example, IT admins can use VMware View to host virtualized user desktops on a central server, then let users access those desktops using secure thin clients. It's not a perfect solution for every case, but increasing network bandwidth and Internet bandwidth, as well as rapid advancements in virtualization technology, could make this model an effective alternative to so-called "fat clients" (i.e., user PCs that are very capable individually but are expensive for IT pros to maintain and manage).

According to a recent Gartner report (G00159622, www.gartner.com), the total cost of ownership

(TCO) of a server-based computing platform that delivers all applications to users is around 50 percent less than unmanaged desktop deployments, and 11 to 18 percent less than well-managed client PC deployments.

Renewed interest in a virtualization-powered thin-client computing model has buoyed the fortunes of thin-client provider Wyse, which has seen its revenue increase. "[The] best barometer here is that while PC company revenue has fallen [in Q4 2008], Wyse's revenue has continued to grow," says Jeff McNaught, Chief Marketing and Strategy Officer for Wyse Technology. "We believe this is an indicator that companies that decided on thin clients to reduce TCO are going ahead with those installations, and not delaying." McNaught points out that the majority of Wyse's thin-client customers are using Microsoft Terminal Services and/or Citrix XenApp with their thin clients, although the company is seeing more customers using Citrix XenDesktop and VMware View recently. If you're using Citrix XenDesktop or VMware View to create hosted virtual desktops (and minimize your rich client overhead), drop me an email—I'd love to hear how cost effective those solutions have been for you.

Consolidating Identity and Security

Consolidation can work wonders in the context of virtualization, but it can reap benefits in other areas of your IT infrastructure as well. Microsoft believes that identity and security have a brighter future together than they do apart, so the company has started to merge the product groups responsible for identity and security into one cohesive unit, dubbed the Identity and Security Business Group. I recently spoke with Microsoft's general manager for that group, Doug Leland, to get a look at what Microsoft has planned for this new division. You can see this interview on page 19.



InstantDoc ID 101791

JEFF JAMES (jjames@windowsitpro.com) is Editor-in-Chief, Web Content Strategist for Penton Media's IT Publishing Group. He specializes in server operating systems, systems management, and server virtualization.

Talk Back

We're always eager to hear reader feedback, so please let us know what's on your mind. Send me an email at jjames@windowsitpro.com, follow me on Twitter @jeffjames3, or give me a call at 970-203-2775. To participate in an online survey about using virtualization as a cost-saving tool, go to <http://tinyurl.com/ccvs5u>.

■ AD Audit Script
■ Cloud Computing

■ A Mobile Future
■ Email Retention

LETTERS@WINDOWSITPRO.COM

AD Audit Genius

Jim Turner's AD audit script ("Track Active Directory Changes," February 2009, InstantDoc ID 100428) is pure genius. I didn't use the script exactly the way the article describes, but the ideas and concepts that Turner presents are worth 100 pages of articles I've waded through elsewhere.

—Marc Casillo

Use the Right Tool

Cloud computing is just one of many tools an organization can use to handle its IT needs. Too many people think of the cloud offerings as all-or-nothing propositions. But it's not necessary or practical to rip out the data center and move every last bit of IT into the cloud. As Jeff James reasons in his IT Pro Perspective column, "Cloud Computing" (January 2009, InstantDoc ID 100943), cloud computing might not be a good fit for large multinational organizations that are subject to strict data regulation. For smaller organizations without a formal IT staff, hosted Exchange services can be a godsend.

We're a small manufacturing company, and if we hadn't already invested in our network hardware and software, cloud services would make a lot of sense for us. We have only about 30 users. We don't use Exchange to schedule meetings, and we don't use public folders. We use Office, but only minimally (we used Office 97 until about 4 years ago), and we rarely share documents or collaborate with customers or vendors. There's really no reason we couldn't use hosted Exchange services. For our Office needs, three or four people would need a local copy of Office. But the rest of us could use either a hosted Office service or Google Docs.

Businesses need to do what makes sense given their level of expertise and willingness to devote the necessary resources to a system. Most companies don't host their own

public website inhouse specifically because web hosting is something that lots of other providers do very well and very cheaply. Unless you have some specific need, why would you do it yourself? Move the stuff that makes sense into the cloud and keep the stuff inhouse that needs to be kept inhouse. As they say, "Use the right tool for the job."

—Peter Diamond

A Mobile Future

In his IT Pro Perspective column, "A Mobile Future" (February 2009, InstantDoc ID 101134), Jeff James wonders how many users have asked about integrating their iPhone with the corporate IT infrastructure. In my experience, the answer is, "A lot." I work in health care, and I've told our hospital's Verizon sales rep to say "No" to BlackBerry and iPhone devices. In fact, the only phone that we support is the powerful, IT-ready HTC Touch Pro. Here's why:

1. Email integration—The phone integrates with our email system, with the help of Verizon's free Wireless Sync service. I don't have to maintain a BlackBerry server, and I don't even have to muck around with our Exchange Server 2003 system to get everything to work for users.

2. Citrix usability—We're a heavy Citrix shop, and Citrix doesn't natively work on BlackBerry devices (you have to get a third-party client). The HTC Touch Pro gives VGA resolution so that doctors can securely access enterprise apps, and it's the only phone I've seen so far that can do it well.

3. General usability—The 3D Flow touch interface is very much like the iPhone's display.

4. Applications—The .NET Compact Framework is maturing, with many free apps available. And we can develop on that platform.

5. Microsoft Office document use—Let's be honest, we don't actually make

ONLINE

windowsitpro.com

Email Retention

B. K. Winstead's "Establishing an Email Retention Policy: The Legal Perspective" (March 5, 2009, InstantDoc ID 101646) is a great article. My employer is about to embark on the same process—what we call the "grooming tool" for email retention. We're going with 90 days or older for deletion out of a user's mailbox. Anything to be kept longer than 90 days needs to be stored in our email-archiving solution. After this policy burns in for a while, we'll be deleting email content over two years old from the archive solution. Anything more than two years old that needs to be kept will have to go to our forthcoming Documentum solution.

—Michael Agents

Thanks for the feedback! Sounds like you also have a carefully considered policy going into effect. I hope you have the same kind of supportive and collaborative team that we've had here at Penton. As I mentioned in the article, the second part of the story is the interview with our IT guys, "Establishing an Email Retention Policy: The IT Perspective" (March 19, 2009, InstantDoc ID 101728).

—Brian Keith Winstead

documents on a cell phone. We read them. We use Office at the hospital, so we can download and read our Office documents with ease.

Corporate policy is easy to maintain with our configuration. If we feel that a phone is compromised, we can change the Active Directory (AD) password for that user. These phones don't even touch our desktop machines; if they did, we'd be worried. Also, we use the free My Mobiler for remote support, if necessary.

Stick with Windows Mobile on a great device with a good carrier that offers small-scale or large-scale sync—and doesn't require much support. Don't write off the Windows Mobile phone!

—Dylan McNeill
InstantDoc ID 101760

Windows IT Pro welcomes feedback about the magazine. Send comments to letters@windowsitpro.com, and include your full name, email address, and daytime phone number. We edit all letters and replies for style, length, and clarity.

ONLINE

windowsitpro.com

Overcome 3 Major BI Optimization Barriers

Learn the most important factors to consider when optimizing your business intelligence (BI) with this Essential Guide from IT powerhouses Microsoft and Intel. You'll find out how to move forward with BI optimization when your data is dispersed across multiple source systems, the data quality in the source systems is poor, and the relational databases are unsuccessful in running your analytical queries.

windowsitpro.com/go/OptimizeBI

Is Windows 7 Right For You?

Join veteran Windows watcher and *Windows IT Pro* Senior Contributing Editor Mark Minasi on May 28 for a clear, comprehensive, independent, and often entertaining look at what Windows 7 can (and can't) do for you. Mark explains what's new in Windows 7 from soup to nuts to save you time and help you make an informed "upgrade or not?" decision.

windowsitpro.com/go/GettingReadyforWindows7

Take Control of Your Email

Learn more about the business issues associated with email storage management and discover approaches for managing this storage while accommodating the needs of the wider business. This white paper includes email archiving solutions designed to help you control and manage your organization's email. This not only helps you address compliance requirements, but also reduces overall email storage utilization, assists with mailbox management, and increases the performance of the email server. You'll also learn how ScriptLogic's Archive Manager can help. Archive Manager captures, indexes, searches, and archives Exchange messages and attachments.

windowsitpro.com/go/TakeControlOfEmail

Humphries

The missing link to
IT resources



Make Social Networking Work for You

Widen your horizon and opportunities with *Windows IT Pro*

I'm usually the last person to jump on a bandwagon, for fear of hurting my pride if I fall off. So I wasn't the first to sign up for social networking. I thought that it would just be a virtual hangout for high schoolers—or for those stuck in the high school frame of mind. But now I realize that social networking can be a great source for staying competitive in your field, a hotline for quick tips and instant solutions, and a collection of contacts. Now you can do all of that with help from the *Windows IT Pro* network you trust.

Twitter

When my coworker asked me if I Twittered, I wasn't sure if I should be offended or flattered. I came to find out that it wasn't some offhand remark; Twitter is another way to stay in touch with your network.

SavvyAsst. Serving as an extension of my monthly column that you love, this account will inform you of helpful resources, free tools, new events, and industry happenings. Come follow me at twitter.com/SavvyAsst.

EmpowerITDev. Empower IT/Dev is the User Group program for the *Windows IT Pro* network. The feed is used to inform user group leaders and members about our User Group (UG) sponsorship opportunities, and to cross-promote meetings, activities, and news for groups that we follow (and ideally partner with). Empower your user group at twitter.com/EmpowerITDev.

LeftBrainStore. Left-Brain.com is the new online resource superstore stocked with educational, training, and career-development materials concentrated on the needs of IT professionals like you. Sign up to receive new product alerts and special Twitter discounts on the resources you

need. Get on the right side of the IT curve at twitter.com/Left-BrainInformationStore.

IT Job Hound. IT Job Hound is an online job-search engine that concentrates on the IT industry. (To learn more, see a previous Your Savvy Assistant post at InstantDoc ID 100047.) This feed will keep you tracking new positions, jobseekers, and trends in your fields of interest. Sniff out the right career or employee at twitter.com/ITJobHound.

Windows IT Pro. As the Twitter feed of the industry's independent resource for product news, information, and community, this account shows how we're in IT with you: twitter.com/WindowsITPro.

LinkedIn

LinkedIn is the most professionally formal site of the social networking world. Treat this as an extension of your resume, where you can record your work history, link up with current and former coworkers and managers, and even compile all of those shining recommendations. Link up with *Windows IT Pro* for your career connections.

Facebook

From what I've heard from IT pros, Facebook is used more for peer development than for career development. But for many professionals out there, this can be a more comfortable and casual forum for your networking needs. *Windows IT Pro* has a Facebook account at <http://tinyurl.com/d5bquf>. Check us out to connect with other IT pros and our experts.

To learn more about IT industry social networking (beyond the cram session at TechEd), contact me at Christan.Humphries@penton.com.

InstantDoc ID 101679



"Windows 7 is very clearly 'Windows Vista done right,' and it looks like those who elected to skip Vista made the right decision."

What You Need to Know About Windows 7 RC

After the delays we saw with Windows Vista, the Windows 7 development cycle has been straightforward and, dare I say, speedy. Microsoft issued a pre-beta version to external testers in October 2008, then shipped a public beta in January. The Windows 7 release candidate (RC) offers a last chance to evaluate this OS. Here's what you need to know about the Windows 7 release candidate.

Where the Release Candidate Fits

In this final prerelease milestone for Windows 7, Microsoft will change only the things that prevent the OS from working properly under certain conditions, the so-called "showstopper" bugs. During the months-long gap between the release candidate and the final release, as Microsoft's partners and customers get ready for the general availability of the OS, Windows 7 will be effectively locked down and "very few changes" will be made to the code, according to Microsoft.

Changes Since the Beta

Microsoft called the Windows 7 beta "feature complete," representing, largely, Windows 7's final form. But an onslaught of feedback after the public beta release prompted changes running the gamut from major functional updates to minor UI tweaks. Some of the interesting ones in the release candidate include the following:

User Account Control changes. Microsoft toned down UAC's constant prompting of users, but external testers weren't happy with some UAC functionality. Now, UAC will run as a high integrity process, ensuring that users are validated via a UAC prompt before they can make changes. Users can alter UAC's security level after validating against a UAC prompt.

User experience changes. Based on user feedback, the new Aero Peek effect is now an option in the Windows Flip (Alt+Tab) pop-up window. The Windows Key offers new keyboard shortcuts aimed at power users. "Needy" applications that prompt users with flashing taskbar buttons are more visually distinct. Taskbar scaling has been improved to display more icons at a time. And Windows 7's new themes support is easier to use and less likely to lose a user's changes.

Windows Explorer. The Windows 7 shell offers a tweaked UI, more obvious drag and drop in the new view styles, support for local (i.e., fixed) FAT32 disks, and a wide range of new icon view arrangements.

Windows Touch improvements. One of the big enhancements to Windows 7 is its globally-available touch interface called Windows

Touch, which has been augmented with an Aero Peek touch gesture and Show Desktop support. The onscreen touch keyboard also now supports multi-touch, allowing such key presses as Ctrl+C and Shift+[letter] for capitalization. And a new multi-touch gesture offers more natural right-click support.

Windows Media changes. While many feel that Microsoft should have used its Zune software in Windows 7 instead of Windows Media Player, the new WMP version in Windows 7 does provide improved Internet radio playback, a cleaner Now Playing window, better power-management awareness, simpler device sync, and custom Jump List improvements.


Internet Explorer 8.0 is now removable. Responding to a potential legal threat from European Union (EU) antitrust regulators, Microsoft made IE 8.0 removable via the standard Windows Features UI. Whether or not this functionality will be included in all versions of Windows 7 is unclear for now.

Control Panel updates. IT pros and admins can lock a Windows 7 PC without first requiring a screensaver. And the High Performance power management scheme is now visible, as it was in Vista.

Hardware support changes. A new Device Stage UI discovers and surfaces the functionality provided by a wide range of devices, printers, and other hardware. The release candidate supports a wider range of hardware than did the beta.

Performance. Microsoft tweaked the performance of this already surprisingly limber OS. It boots up, runs, sleeps, resumes, and shuts down quicker than Vista and the beta release.

Recommendations

Windows 7 is very clearly "Windows Vista done right," and it looks like those who elected to skip Vista made the right decision. Although there are few advantages to migrating from Windows XP to Windows 7 (compared with migrating from XP to Vista), and even fewer for those upgrading from Vista to Windows 7, XP-based environments should move to Windows 7 as soon as possible because of its usability, manageability, and security improvements. This OS is the most feature-packed and secure version of Windows yet, and if the RC version is any indication—and it is—Microsoft has a winner on its hands. 

InstantDoc ID 101640

PAUL THURROTT (thurrott@windowsitpro.com) is the news editor for *Windows IT Pro*. He writes a weekly editorial for *Windows IT Pro UPDATE* (www.windowsitpro.com/email) and a daily Windows news and information newsletter called *WinInfo Daily UPDATE* (www.wininformant.com).

Visit Sunbelt Software at Microsoft Tech Ed - Booth # 111

Kiss your antivirus bloatware goodbye

Special
Competitive
Upgrade Price:
\$10 per seat!



VIPRE ENTERPRISE

TEST DRIVE

Next Generation of Total Malware Protection

Until now, antivirus engines have been Franksteins, bolted together from bits and pieces of different products. They're slow, full of bugs, and hard to manage.

VIPRE Enterprise is a revolutionary new approach. It's built from scratch as the all-in-one antivirus, antispysware, anti-rootkit solution that gives you complete endpoint malware protection **without hogging resources!** It's fast, powerful, and easy.

Plus, advanced anti-malware technology protects your system against the new wave of malware threats. No more juggling multiple programs. No more dealing with user complaints about slow workstation performance.

- **COMPLETE!** All-in-one protection from today's malware.
- **FAST!** High-performance and low impact on system resources.
- **EASY!** Manage everything easily from one command screen.
- **RELIABLE!** Configurable, real-time monitoring technology.
- **AFFORDABLE!** Low \$10 per seat pricing to save you money.


Why struggle with slow resource hogs when you can manage ALL your malware threats with one fast, easy application?

Curious? Download your FREE copy of VIPRE Enterprise and give it a test drive.

When you compare VIPRE Enterprise to Symantec, McAfee, Trend Micro or whatever antivirus program you're using, **you WILL want to switch!** Don't worry, though. You can get VIPRE Enterprise at our competitive upgrade price of **only \$10 per seat!**



The configurable Command Center puts all the information you need in one place. Manage individual agents, quarantines, threats, and more.



How does your current software compare? VIPRE Enterprise scans at a brisk 13.95 MB/sec and uses just 27% of CPU and 50 MB of RAM. In idle, it uses a mere 13.3 MB RAM with a disk footprint of just 113 MB. You'll hardly notice it's running!



Sunbelt Software

Download VIPRE Enterprise today and get your own home version of VIPRE to keep FREE as our gift to you!

Download now: **www.TestDriveVipre.com**

Sunbelt Software Tel: 1-888-688-8457 or 1-727-562-0101 Fax: 1-727-562-5199 www.SunbeltSoftware.com sales@sunbeltsoftware.com

© 2009 Sunbelt Software. All rights reserved. VIPRE Enterprise is a trademark of Sunbelt Software. All trademarks used are owned by their respective owners.

New licenses are available for \$10/seat up to 500 seats, minimum 10 seats. For customers with over 500 seats, please call for special pricing. Available for a limited time and subject to change without notice. See website for more details.



"The ability to determine how many processor cores Windows uses can be quite useful for smoking out occasional lockups."

2 Useful Bcdedit Options

Take control of DEP and the number of processor cores

In "Bcdedit Basics" (March 2009, InstantDoc ID 101168) and "Booting Up with Bcdedit" (www.windowsitpro.com, InstantDoc ID 101362), I showed you how to use Bcdedit to control boot options in Windows Vista and later. (Recent versions of Windows lack the boot.ini text file that pre-Vista systems offered to control boot options.)

Recall that Vista and later keeps its boot information in the Boot Configuration Database (BCD). Your BCD can contain more than one OS entry—a set of configuration information that BCD can use to boot a particular OS on your computer. Most of us have only one OS on our computer, but we might still want more than one OS entry so that we can boot our systems with various options for debugging or analysis. Recall also that Bcdedit identifies OS entries not by user-friendly names but instead by random GUIDs such as {9c219fb1-bb55-11dd-97ac-804080387aa6}. Thus, before you can add or subtract options from a particular OS entry, you'll need that OS entry's GUID. Now, I'll put all that background to good use by showing you how to benefit from two of Bcdedit's boot options.

Enabling/Disabling DEP

One useful OS boot option is the *nx* entry, which enables or disables Windows' Data Execution Prevention (DEP) security feature. By default, Windows enables DEP, which constantly watches for worms attempting to take control of the system. If DEP thinks your OS or application is under attack, it shuts down the affected software. DEP is a good idea and is a major contributor to the fact that we haven't seen a widespread Windows worm since late 2003, but it can burn up a lot of CPU cycles and slow your system down noticeably. In my opinion, DEP's worm-fighting value far exceeds that lost value in system speed, but I do recommend disabling DEP in two cases: on test systems or systems on networks that aren't connected to anything. Many organizations' test systems are either hand-me-downs or virtual machines (VMs), neither of which are very speedy. To disable DEP on those systems, open an elevated command prompt and type

```
bcdedit /set [guid] nx AlwaysOff
```

Notice that Bcdedit doesn't require a GUID; if you skip it, Bcdedit assumes you want it to work on the currently active OS entry. And please be extremely careful when editing your BCD: You could easily render your system non-bootable. As I've suggested in previous articles, create a separate OS entry and test options there. That way,

if you end up disastrously goofing up, you've always got your basic OS entry to fall back on. To re-enable full DEP, type

```
bcdedit /set nx AlwaysOn
```

Determining Number of Processor Cores

Another useful Bcdedit option is the *numproc* option, which lets you determine how many processor cores Windows uses. For example, my laptop runs two processor cores, but I recently instructed Windows to use only one processor core by typing

```
bcdedit /set numproc 1
```

After I rebooted, I opened Task Manager to see that Windows was running on one core. Later, I found this option useful for smoking out occasional lockups. Sometimes, I've seen applications that crash or lock up mysteriously on one system but not on others—only to realize that the difference was the number of processors in the different systems. Creating a separate OS entry and configuring that OS entry to run on just one processor (or perhaps two, three, or more) is an easy way to test whether a certain app is experiencing a multiprocessor problem. I like the flexibility of being able to specify the number of processors, but if you need only to compare single-processor scenarios with multi-processor scenarios, you might run across an alternative Bcdedit setting called *onecpu*, which takes the parameters *true* or *false*. Which means that the command

```
bcdedit /set onecpu true
```

has the same effect as

```
bcdedit /set numproc 1
```

And before you ask, I have no idea why Microsoft offers such redundant options.

More Useful Than You Think

Put *nx* and *numproc* in your tweaking and troubleshooting toolkit. You might think they're minor tools, but I bet you find them to be more useful than you thought—I know I did.



InstantDoc ID 101580

MARK MINASI (www.minasi.com/gethelp) is a senior contributing editor for *Windows IT Pro*, an MCSE, and the author of 25 books, including *Administering Windows Vista Security: The Big Surprises* (Sybex). He writes and speaks around the world about Windows networking.

Windows[®] IT Pro

The Impact of Disk Fragmentation on Servers

By David Chernicoff

Published: May 2009



The Impact of Disk Fragmentation on Servers

Testing Server Disk Defragmentation

IT professionals responsible for server hardware well understand the value that professional grade disk defragmentation software brings to their servers. Storage servers can experience high levels of disk thrashing (the constant writing and rewriting of small amounts of data) caused from excessive file fragmentation.

Problems in delivering services to users however are difficult to directly trace to server fragmentation issues. Network and application issues have a much more visible impact on the performance of network-based services, especially when problems with those functions are encountered. But with the negative impact on ROI that network performance problems cause, IT pros would be ill-advised to overlook the advantages that assuring the optimization of the underlying hardware infrastructure can bring. Optimal disk performance translates into better ROI. Testing will bear this out.

We tested the impact of server disk defragmentation by looking at common tasks that network servers, both physical and virtual, encounter, ranging from maintenance tasks such as server backup and anti-virus scans, to basic knowledge worker tasks involving opening files stored on the host server and virtual machines, and manipulating email. We also looked at tasks that are more taxing on the server, such as database queries, index creation, and bulk updates. Each test was performed as the sole task on the server.

When considering the results of our testing keep in mind that a production environment will see significantly heavier server use, which results in much greater potential for ongoing disk fragmentation. In your production environment with dozens, if not hundreds, of users touching your server storage simultaneously, your disk fragmentation can become severe in a very short time. Preventing this fragmentation from affecting server performance is an ongoing process.

→ Contents

Testing Server Disk Defragmentation	2
The Testing Environment.....	3
The Tests.....	4
File Copy.....	4
Document Open	4
Backup.....	5
Anti-Virus Scan	5
VHD Start	5
VHD Save	6
Server Application Tests	6
Exchange Test One	6
Exchange Test Two.....	7
SQL Server Bulk Insert.....	7
Table Key Creation	7
SQL Query 1	7
SQL Query 2	8
Conclusion	8



The Testing Environment

For our benchmark tests we used an HP ProLiant DL380 G5 equipped with dual quad-core 2.83 GHz Xeon processors, each with a 2x6MB L2 cache, 16 GB of RAM and seven 72 GB 10,000 RPM SCSI drives attached to an HP Smart Array P400 controller that has a 256 MB cache and that supports both serial-attached SCSI and SATA drives. The volumes we tested against were 30 GB, 80 GB, and 175 GB. We used a 500 GB 7200 RPM locally attached SATA drive for backup only. The server operating system was Microsoft Windows Server 2008 Enterprise; the application server software installed in VHDs was Microsoft SQL Server 2008 and Microsoft Exchange Server 2007. All server software was updated with service packs, patches and hotfixes current as of February 2009. The disk defragmentation software used was Diskeeper Server.

The seven SCSI drives attached to the array controller were configured as two physical drives. We used the first physical drive, comprised of two drives configured as a RAID 0 stripe set for maximum performance, for the installation of the operating system and all related files. We configured the remaining five drives as a RAID 5 stripe set to be representative of the type of hardware storage configuration found in most business environments. We performed all applications, VHDs, and tests on the RAID 5 stripe set. The volume size was dependent upon the test level.

As an example of the effect fragmentation can have, the screen capture in Figure 1 shows the Diskeeper fragmentation analysis of a severely fragmented disk. The severe fragmentation documented here will have a negative impact on storage performance.

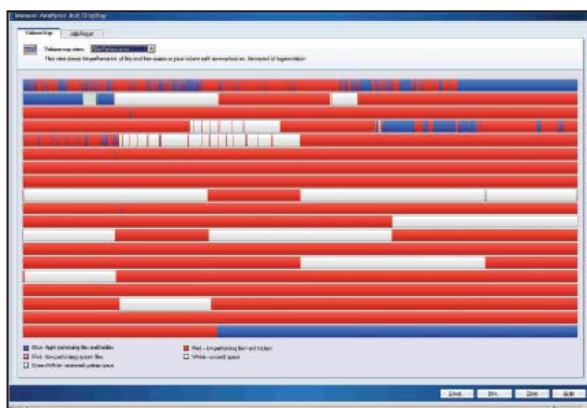


Figure 1: Fragmentation map of a heavily fragmented disk

We tested three levels of fragmentation, described herein as low, medium, and high. We used the Diskeeper Diskcrusher fragmentation utility to create fragmented files and directories. We ran all tests a minimum of three times with the results reported here being the average of all test runs.

	Low	Medium	High
Number of files	101,652	1,220,660	2,087,158
Avg. Number of Fragments per File	3.21	1.69	2.30
Number of Fragmented Files	99,074	613,221	1,994,117
Number of Excess Fragments	225,216	840,076	3,005,400
Percent Fragmented – Volume	40%	50%	84%
Percent Fragmented – Data	51%	58%	91%
Free Space	22%	15%	15%

Table 1: Fragmented disk test configurations

As shown in Table 1 the level of fragmentation and the number of affected files increases with each testing tier. The level of fragmentation you'll encounter in production environments is dependent upon the level of use and types of applications the server deals with. In all likelihood, if your server storage levels are consistently exceeding 75 percent or so, you've begun aging data off of the servers or you're planning to add additional storage. While fragmentation isn't a direct result of reduced capacity, the chances for fragmentation increase as free storage space decreases and the operating system is forced to write data into an ever-increasing number of non-contiguous spaces.

By using an automated defragmentation process, the same disk volume sees absolutely minimal fragmentation even though it is in continual use by applications and users (Figure 2).




Figure 2: Fragmentation map after automated defragmentation by Diskeeper.

We ran each set of tests for three iterations, and then defragmented the storage using Diskeeper to reduce or eliminate the disk fragmentation. We repeated each test (also for three iterations) and averaged the results. In the following test descriptions and analysis, the comparisons are all before and after defragmentation at each specific fragmentation level tier. We did not do cross-tier comparisons. All test times are reported in seconds.

	Low	Medium	High
Number of files	101,652	1,220,660	2,087,158
Percent Fragmented – Volume	0	0	0
Percent Fragmented – Data	0	0	0
Avg. Number of Fragments per File	0	0	0
Number of Fragmented Files	0	1	1
Number of Excess Fragments	0	2	4
Free Space	22%	15%	15%

Table 2: State of fragmentation after Diskeeper has been run

The Tests

In our first set of tests we look at common server tasks that are likely to be affected by disk fragmentation. These tasks are all primarily storage related; that is, the performance of the storage media will have a primary impact on the performance of these tasks.

File Copy

In the file copy test, a folder containing 5 GB worth of files and sub-directories was copied from the test volume to the boot volume of the server. To minimize variables, the copy was done locally, not across the network. We timed the test using a stopwatch. This is one of the most basic tasks done with server data and, in a severely fragmented environment, showed some of the most significant performance improvements.

File Copy Tests (measured in seconds)

Low – Fragmented	44
Low – Defragmented	39
Medium – Fragmented	72
Medium – Defragmented	60
High – Fragmented	97
High – Defragmented	54

The basic task of moving data from one location to another on the server shows that a fragmented disk has a major negative impact on the file copy. Even the lightly fragmented low-level test showed an improvement in copy time of over 11 percent, while the copy that was done from the very highly fragmented drive improved in time by almost 45 percent. Given how common the file copying task is the benefit is clear. Defragmented disks are a significant time saver for common user tasks.

While the limiting factor in doing a file copy from the server to the client might be the available network bandwidth, as technologies such as Gigabit Ethernet become more common, the base limiting factor will be how fast the operating system can feed data to the network request, which is directly impacted by fragmentation of the data on the local drive.

Document Open

In this test, a 100-page Microsoft Word document was opened from the server to a Windows XP client running Microsoft Office 2007. The size of the document was 3.3 MB.

Document Open Tests (measured in seconds)

Low – Fragmented	11.7
Low – Defragmented	10
Medium – Fragmented	12.7
Medium – Defragmented	10.7
High – Fragmented	14.7
High – Defragmented	10.3

Our test results showed performance improvements of upwards of 30 percent. In the case of any file load from server to client the performance improvement will be determined by just how badly fragmented is the file located on the server. In our tests, the file was clearly badly fragmented, significantly so at the highest level of fragmentation testing. To prevent this type of file fragmentation, the best methodology is an ongoing background file defragmentation process, the benefits of which are clearly demonstrated by this test. And given how often this type of task is performed in most business environments, the value of the defragmentation cannot be understated. As shown in this and the File Copy test, basic data manipulation is much faster on defragmented storage.

Backup

In the first test, we backed up the test volume using disk-to-disk backup as supported by Windows Server Backup, which is a component of Windows Server 2008. Backup was done using the VSS copy method, which is designed to work with other backup tools that would require that the archive and backup information in the files remain unmodified. We backed up to a SATA-attached dedicated hard drive that was reformatted between tests. Timing was done using the backup application.

Backup Tests (measured in seconds)

Low – Fragmented	1193
Low – Defragmented	1130
Medium – Fragmented	2787
Medium – Defragmented	2300
High – Fragmented	6960
High – Defragmented	6620

While different backup tools will be differently affected by disk fragmentation, our tests showed one simple fact; defragmented disks back up faster. Individual runs demonstrated performance improvements of up to 20 percent with our test data set and the built-in Windows Server backup. Our least effective test result, a large data backup that can represent a significant amount of time, still showed an improvement of 5 percent. Our highest report results, which averaged a 17 percent reduction in backup time, shows that reducing or eliminating disk fragmentation prior to backup will allow larger amounts of data to be backed up, especially if time is a constraint in your backup process. If backup is run as a background application, reduced fragmentation will allow for lower resource consumption necessary for the backup process, minimizing further the impact of the backup on active users of the storage.

The single, consistent result that appears in all of our tests is that defragmented server drives using Diskeeper deliver better performance.

Anti-Virus Scan

For the AV scan test, we performed a complete scan of the test volume using the Kapersky Lab AntiVirus Version 6 Windows Server software, current as of February 2009. The default configuration of the AV software was used with only the test volume selected for scanning. Timing was done using the AV application.

Anti-Virus Scan Tests (measured in seconds)

Low – Fragmented	256
Low – Defragmented	238
Medium – Fragmented	1485
Medium – Defragmented	1359
High – Fragmented	4428
High – Defragmented	4004

Many factors will have an impact on the speed of a complete anti-virus scan of your storage. The way the scanner works, the total number of files that need to be scanned, the size of the files, and the fragmentation level of the disk all have a direct impact on the length of the AV scan process. In our tests with the Kapersky Lab AV solution, the disk defragmentation resulted in upwards of a 10 percent performance improvement—with the improvement being more significant as the test drives increased in size, number of test files, and fragmentation.

VHD Start

This test measured the amount of time it took to launch the saved test virtual machine. The VM was launched from a saved state and timing stopped when the Hypervisor manager reported that the VM was successfully started.

VHD Start Tests (measured in seconds)

Low – Fragmented	62.3
Low – Defragmented	51
Medium – Fragmented	60.7
Medium – Defragmented	58
High – Fragmented	55.3
High – Defragmented	47

With as much as a 17 percent improvement in the start time of the test virtual machine, the effects of fragmentation on the VHD are clear. This fragmentation will also impact the performance of the VM itself, because all of the additional I/O necessary to read from a severely fragmented VHD will reduce the performance of the virtual computing environment. Fragmentation must also be watched if your VMs are configured with the dynamic disk option, which allows the virtual machine to grow the size of its storage as necessary. This means that as the size of the VHD grows it will continue to fragment into the available space on the hard drive. Making sure that the host machine hard disk is regularly defragmented and managed will improve the performance of virtual machines running on the host and allow for the use of dynamic disk allocation within the VM without danger of disk performance issues.

Even with significant free space of the disk, as shown by the white space in the fragmentation map (Figure 3), major fragmentation can still occur even without VHD test volume.

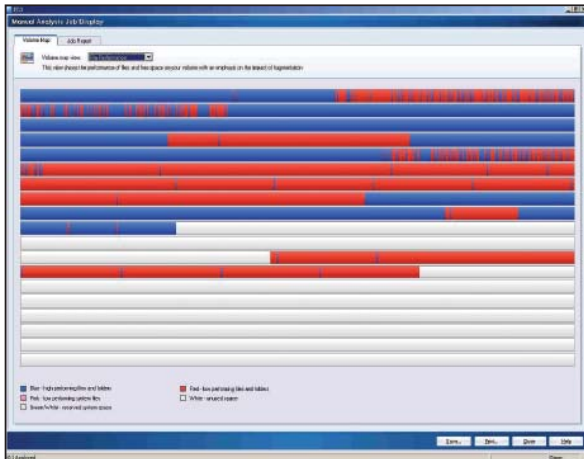


Figure 3: Fragmentation map of VHD volume

VHD Save

This test measured the length of time required to save the test virtual machine. From the Hypervisor manager, the running machine was saved and timing stopped when the manager reported the save complete.

VHD Save Tests (measured in seconds)

Low – Fragmented	365.3
Low – Defragmented	271.7
Medium – Fragmented	409.3
Medium – Defragmented	402
High – Fragmented	447.7
High – Defragmented	390.3

With test results indicating as much as a 25 percent performance improvement after defragmentation, the VHD Save tests show quite clearly the effect of writing a very large file to a fragmented hard drive. The more fragments on the drive the less likely it will be that a large file can be written contiguously. And in the world of virtualization, large files are the standard, and the need to be able to read and write those files with a minimum of fragmentation is a requirement to meet the basic ROI needs of the enterprise.

Automated background defragmentation results in a major reduction in fragmentation even with an active VHD (Figure 4). Regular use of the background defragmenter will continue to minimize fragmentation.

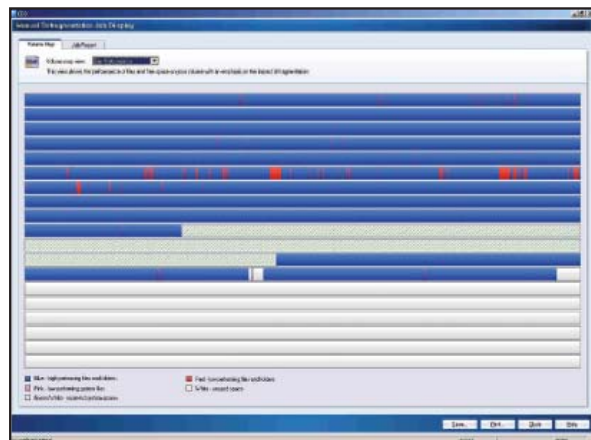


Figure 4: Fragmentation map after automated defragmentation by Diskeeper.

Server Application Tests

In the server application tests we looked at the impact of fragmented storage on server-based applications. Other factors will have an impact on the overall performance of these applications; optimizing storage strategies, including defragmentation, reduces the impact of storage performance on the overall application performance.

Exchange Test One

In this first Exchange test, the client, a Windows XP Professional Workstation running Office 2007, uses Outlook to open 100 messages from the server. One hundred messages are highlighted then opened simultaneously. Timing starts when the open is launched and stops when all of the messages have been opened and console control returns.

Exchange Test One (measured in seconds)

Low – Fragmented	7.7
Low – Defragmented	7
Medium – Fragmented	10.7
Medium – Defragmented	8.6
High – Fragmented	18.4
High – Defragmented	11.6

While the impact of server fragmentation gets significantly greater as the disk becomes more fragmented, even the common lower levels of fragmentation will have a large impact on user response time when you consider that hundreds of users may be accessing the data store at the same time. Delayed response time for email users is a generator of a large percentage of help desk calls, and implementing a defragmentation strategy can help to solve the problem. As our tests show, allowing the data to become seriously fragmented can have a major negative impact on the

Exchange user experience with a 40 percent reduction in performance in our highly fragmented test environment. Good defragmentation strategies result in fewer help desk calls.

Exchange Test Two

In this test, the contents of an existing folder were moved to a new folder. Time to complete was measured from the client side.

Exchange Test Two (measured in seconds)

Low – Fragmented	9
Low – Defragmented	8
Medium – Fragmented	13.8
Medium – Defragmented	9
High – Fragmented	24.9
High – Defragmented	12.3

A new folder was created and the contents of the Inbox were moved to the new folder. With our heavily fragmented test environment showing a greater than 50 percent performance improvement after defragmentation it's clear that this test was extremely sensitive to higher levels of fragmentation on the server. If users are often found reorganizing the data in the Exchange mailbox, the impact of fragmentation can be quite severe.

SQL Server Bulk Insert

We tested SQL Server 2008 with a bulk insert of 50,000 rows of data. The bulk insert is often the fastest method of getting data into a SQL Server database.

SQL Server Bulk Insert Tests (measured in seconds)

Low – Fragmented	22.1
Low – Defragmented	20.9
Medium – Fragmented	31
Medium – Defragmented	25
High – Fragmented	53.3
High – Defragmented	33.4

As has been seen with the Exchange tests, a highly fragmented database structure can have a severe negative impact on loading and extracting data from server applications, with our test showing a performance improvement of 40 percent in the most heavily fragmented environment. Because Microsoft offers APIs for moving open files, defragmentation software is able to safely work on database files without risk of data loss or corruption. Loading data into a defrag-

mented environment not only improves load times but reduces the amount of disk thrashing necessary to manipulate the data and the amount of work that is necessary to later defragment the database.

Table Key Creation (measured in seconds)

	Table 1	Table 2
Low – Fragmented	12.5	15.9
Low – Defragmented	12	14.9
Medium – Fragmented	14.1	18.23
Medium – Defragmented	12.4	17.1
High – Fragmented	25.5	32.4
High – Defragmented	20.6	25.3
	Table 3	Table 4
Low – Fragmented	26	35.4
Low – Defragmented	24.2	33
Medium – Fragmented	32.3	49.1
Medium – Defragmented	30.4	43.8
High – Fragmented	51	68.8
High – Defragmented	46.7	61.3

In this test each table was opened, a field was selected as the primary key, and the change was saved. The table key creation times are directly related to how much data SQL Server had to touch, and the level of fragmentation that had to be dealt with. SQL Server 2008 does a very good job of managing its databases, but defragmentation shows appreciable improvement in the performance of tasks such as this with a performance improvement of over 11 percent in the most fragmented environments.

With the SQL queries, the two tests differ primarily in the amount of data that SQL Server returns in response to the query. The tests depict the effects of manipulating the data on a fragmented drive with peak performance improvements of approximately 18 percent.

SQL Query 1 – Simple (measured in seconds)

Low – Fragmented	23.9
Low – Defragmented	22.3
Medium – Fragmented	28.2
Medium – Defragmented	24.8
High – Fragmented	43.5
High – Defragmented	33

SQL Query 2 – Complex (measured in seconds)

Low – Fragmented	35.3
Low – Defragmented	33.3
Medium – Fragmented	41.5
Medium – Defragmented	38.5
High – Fragmented	61.3
High – Defragmented	50.8

Conclusion

The single, consistent result that appears in all of our tests is that defragmented server drives using Diskeeper deliver better performance.

Every application that touches the hard drive will benefit from a good tool that defragments and manages the files on your servers.

Almost every role filled by Windows servers in your computing environment will benefit from the use of disk defragmentation software. The simplest file and print services delivery requires a significant amount of disk I/O and will easily benefit from file defragmentation. As our simple tests show, even Exchange and SQL Servers benefit from defragmentation; reading and writing data with either application simply

works better when the files are not fragmented. The result is improved performance.

Throwing more storage resources (hardware) at a problem should be the last resort, because it only masks the potential problems that intelligent disk defragmentation addresses.

Quicker response time in databases and mail servers means that more time is spent getting work done, rather than waiting for information to be delivered.

Diskeeper is the only true server defragmentation software that runs silently in the background, continually improving performance.

With the current economic and business environment, maximizing ROI becomes even more critical. Adding Diskeeper to your server toolkit gives you the ability to get the maximum speed from your storage subsystems of your existing hardware.

Our test results showed performance improvements of upwards of 30 percent.

David Chernicoff is a technology consultant with a focus on the mid-market space, *Windows IT Pro* Senior Contributing Editor, founding Technical Director for *PC Week Labs* (now *eWeek*), former Lab Director for *Windows NT Magazine/Windows 2000 Magazine* (now *Windows IT Pro*) and formerly Chief Technology Officer for a network management tools ISV. David has been writing computer-related

feature and product reviews for more than 20 years and is coauthor of a number of operating system books, ranging from the *Windows NT Workstation: Professional Reference* (New Riders Publishing), to the *Microsoft Windows XP Power Toolkit* (Microsoft Press), as well as over a dozen eBooks on topics ranging from network switching topologies to production FAX technology.


"Perhaps the biggest mark against Vista in the business world is that there isn't any real ROI."



Windows Vista Shortcomings as a Business OS

Cost, speed, compatibility, and other problems are keeping Vista out of businesses

Windows Vista sales numbers are good because Vista comes on nearly all new PCs—which means that these sales numbers are mainly from consumer systems. By now, it's clear that businesses are in no rush to roll out Vista. In fact, the vast majority have eschewed Vista in favor of sticking with Windows XP. Don't get me wrong—I've used Vista since before its official release, and I still use it today. But I was hoping for something better. Here are some of the marks against Windows Vista as a business OS.

- 10 Cost**—Vista just plain costs more than XP. It requires more powerful hardware than XP, usually requiring upgrades to install, but that's just the beginning. Many applications also need to be replaced with Vista-compatible versions, a cost that can really add up in the enterprise.
- 9 Speed**—It's only fair to expect that every Microsoft OS is bigger and slower than the previous version because the company adds more "customer requested features" to spur sales. However, Vista runs slow on all but the fastest systems. It's not so much each individual item but the sum of all the parts that makes Vista appear bloated and sluggish. And despite the hype, SP1 didn't rectify this situation.
- 8 UAC**—Theoretically, User Account Control (UAC) is a good idea, but in all the time I've run Vista, I've never benefited from it—not even once. That's more than two years of daily use. However, UAC has hassled me for confirmation thousands of times. I like security, but the best thing I did for my productivity was to turn off UAC.
- 7 Ctrl+Alt+Del doesn't always work**—Some program incompatibility is expected with a new release, which often means using Task Manager to end your unresponsive programs. However, with Vista, Task Manager no longer comes up reliably when you hit Ctrl+Alt+Del. Ironically, it fails most often when a program is hung.
- 6 The disappearing Map network drive function**—XP's Map Network Drive option is always available on the Windows Explorer interface—exactly where you would use it. With Vista, this option appears only in the Windows Explorer view of Computer.
- 5 Mobile device compatibility**—The move to Vista could forcibly retire many of your mobile devices—and replacing your mobile devices won't be cheap. My iPAQ 3815 that worked fine with XP and ActiveSync won't work with Vista's Sync Center.
- 4 Sleep**—Improving XP's sleep function was a great idea, but it isn't really an improvement if it doesn't work right. Vista's new hybrid sleep is unreliable. I've seen many systems become unresponsive when coming out of sleep mode, and you're often greeted with the new black screen of death or the system hangs on the logon screen.
- 3 Windows Explorer settings**—Like a buried tick, Windows Explorer's refusal to remember your folder settings is one of those little annoyances that grows on you over time. It works about 99 percent of the time, but every now and then Windows Explorer loses the settings for a given folder.
- 2 Wireless manager**—One of Vista's biggest problems is its support for connecting to Wi-Fi networks using the *Connect to a network* dialog box. Even with the latest service pack, Vista often refuses to automatically connect to saved configurations. Finding your secured Wi-Fi network often requires you to repeatedly click the refresh button—it displays a different set of networks on every click.
- 1 ROI**—Perhaps the biggest mark against Vista in the business world is that there isn't any real ROI. Vista makes some tasks easier, but it makes other tasks more difficult. I switched back and forth between Vista and XP for more than a year and saw no real advantage to using Vista and felt no real loss when I used XP. Combine that situation with the added costs of running Vista, and it's no surprise that businesses have stayed away in droves. 

InstantDoc ID 101571

MICHAEL OTEY (motey@windowsitpro.com) is technical director for *Windows IT Pro* and *SQL Server Magazine* and author of *Microsoft SQL Server 2008 New Features* (Osborne/McGraw-Hill).



"Every desktop object requires desktop heap—memory to store UI objects, such as windows and menus."

Conquer Desktop Heap Problems

Diagnose problems that limit the memory available to Windows desktop sessions

As an administrator, you've probably run into a desktop heap issue and know how tough this type of problem is to solve. First you have to identify that your symptom is related to an exhaustion of desktop heap memory, then discover which process or service is consuming the greatest amount of desktop heap, and finally determine what registry parameter to change to solve the problem. Here I'll explain how to quickly identify whether your system is running out of desktop heap. (See the box below for a list of desktop heap problem symptoms.) Then I'll describe tools and best practices you can use to help resolve the problem.

Windows Internals Background

You'll need some Windows internals knowledge to understand how desktop heap problems occur. Windows 2000 and later systems have a configurable area of kernel mode memory called *session space*. Session space represents a user's logon environment—each user's sandbox of windows and desktops.

Each session contains window stations that act as a security boundary for desktops. Although the term "desktop" may recall the interactive desktop that each user sees when logged on to Windows, not every desktop interacts with the user. Each window station contains desktop objects, and every session will have one interactive window station named WinSta0 that users see when they log on to their systems.

Another way to conceptualize the desktop tree system is that every Win32 thread belongs to a desktop. Every desktop belongs to a window station; one window station per session interacts with the user, while the rest do not. And every window station belongs to a session. A typical system might have Session 0 and Session 1. Session 0 is where services typically run and also represents the console (prior to Windows Vista). Any other session, such as Session 1 or Session 2, typically represents a Terminal Services or Fast User Switching session.

What Is Desktop Heap?

Every desktop object requires memory—desktop heap—to store UI objects such as windows and menus. When applications require a UI object, functions within *user32.dll* are called, and desktop heap memory is allo-

cated. There is one desktop heap per desktop, and the heap memory itself is allocated from session-view space, a subset of session space.

While this process of allocating desktop heap memory works behind the scenes, there are two primary scenarios in which failures can occur. Session-view space can become fully utilized so that no new desktops can be created. This scenario can occur when multiple services run under a nonlocal system-specific user account, creating a new desktop for every instance of the service. In the second, more common, scenario, existing desktop heaps can become fully utilized, so that threads running in that desktop can't use more desktop heap memory. This scenario can be caused by running many instances of the same process or by a process that has heavy UI object usage.

Diagnostic Tools

Let's look at ways to make diagnosing desktop heap exhaustion issues easier. Desktop Heap Monitor 8.1 (Dheapmon), available at tinyurl.com/Dheapmon, is useful for Windows XP or Windows Server 2003 systems. The tool provides a user-friendly menu describing the total number of desktops, sessions, and window stations. The output, which Figure 1 shows, displays each desktop's utilization percentage.

The most important numbers are those in the Used Rate (%) column, which will help you determine whether any of the desktops being monitored are becoming fully utilized (90 percent or more). Another key number is Total Desktop—the total amount of memory allocated by all desktops. If this number approaches the total size of the session-view space, Windows can't create any more new desktops in a session. If this happens, you may need to change a registry value to increase the default session-view size. Table 1, page xx shows, default session-view sizes for Win2K, XP, and Windows 2003.

Before changing registry values, you should try to identify the process(es) consuming large amounts of desktop heap so that you

become aware of the conditions on your system causing the depletion of desktop heap memory. One of the easiest ways to identify a desktop heap hog is by using Task Manager: On the Processes Tab,

Possible Symptoms of a Desktop Heap Problem

- Application startup failures (0xc0000142).
- Scheduled tasks fail to launch.
- Processes silently fail to run.
- UI elements fail to redraw properly.
- An event 243 (*A desktop heap allocation failed*) is logged in the system log.

■ WHAT WOULD MICROSOFT SUPPORT DO?

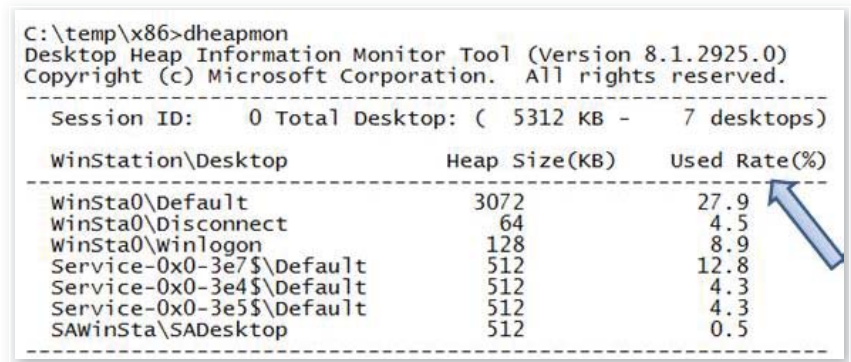


Figure 1: Dheapmon output

Table 1: Default Session-View Sizes

OS	Size if no registry value configured	Default registry value
Windows 2000 *	20MB	None
Windows XP	20MB	48MB
Windows Server 2003	20MB	48MB

* Settings for Windows 2000 are with Terminal Services enabled and hotfix 318942 or a later version of the kernel installed.

click View, Select Columns, and check USER Objects. Click the top of the column to change the sort order to descending, so you can see the application or service consuming the most desktop heap resources. The desktop heap hog may indicate a problem related to the service or application requiring further investigation—and simply adjusting registry settings to work around the problem might only mask the real issue. You can also use the information in Task Manager's USER Objects column to determine which application or service is consuming the largest amount of desktop heap on a Vista or Windows Server 2008 system.

Session-View Space Settings

For XP, Windows 2003, and Win2K, you can configure session-view space size by using the SessionViewSize registry value (REG_DWORD). You specify the size in megabytes. For Vista and later, this value doesn't apply because the session-view space grows as needed. The values in Table 1 are specific to 32-bit x86 systems not booted with the /3GB switch. You must reboot your system to effect this change. You specify the value under the subkey HKEY_LOCAL_MACHINE\SYSTEM\

CurrentControlSet\Control\Session Manager\Memory Management.

If you need to change the size of a specific desktop heap (that is, when Used Rate (%) approaches 90%), you have two possible ways to do so, based on whether one of two conditions exists. The first condition occurs when

Dheapmon data reveals a high used rate for a desktop heap belonging to an interactive window station (WinSta0) and isn't the Disconnect or Winlogon desktop. In this case, you can configure the desktop's heap size using the SharedSection registry value (the second value—3072—for the SharedSection= entry in the registry listing in Figure 2). I'll explain these registry values shortly.

The second condition occurs when the Dheapmon information reveals a high used rate for a desktop heap belonging to a non-interactive window station. In this case, you can also configure the desktop's heap size using the SharedSection registry value (the third value—512—for the SharedSection= entry in Figure 2). The size of each desktop heap allocation is controlled by the registry subkey HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\Session Manager\SubSystems\Windows. The default data for this registry value will look something like that in Figure 2.

```
%SystemRoot%\system32\csrss.exe ObjectDirectory=\Windows
SharedSection=1024,3072,512 Windows=On SubSystemType=Windows
ServerDll=basesrv,1 ServerDll=winsrv:UserServerDllInitialization,3
ServerDll=winsrv:ConServerDllInitialization,2 ProfileControl=Off
MaxRequestThreads=16
```

Figure 2: Sample default data for registry value controlling desktop heap size

As mentioned, the numeric values following SharedSection= control how desktop heap is allocated. These SharedSection values are specified in kilobytes.

- The first SharedSection value (1024) is the shared heap size common to all desktops. This memory isn't a desktop heap allocation, and you should not modify this value to address desktop heap problems.
- The second SharedSection value (3072) is the size of the desktop heap for each desktop associated with an interactive window station (WinSta0), except for the Disconnect and Winlogon desktops.
- The third SharedSection value (512) is the size of the desktop heap for each desktop associated with a non-interactive window station (usually a service). If this value isn't present, the size of the desktop heap for non-interactive window stations will be same as the size specified for interactive window stations.

Vista SP1 and Server 2008 Changes

In 32-bit Vista SP1 and Server 2008, session-view space is now a dynamic kernel address range, and the SessionViewSize registry value is no longer used. This improvement explains in part why you might have fewer desktop heap issues running Vista or Server 2008 than with earlier Windows versions. Also, the second SharedSection value has changed to 12,288KB, the value for interactive desktop heaps.

You're now better equipped to recognize desktop heap problems and resolve such issues on your own. As always, I welcome your questions or stories about your own Windows troubleshooting experiences, with desktop heap or other OS issues.

InstantDoc ID 101701

MICHAEL MORALES (morales@microsoft.com) is a senior escalation engineer for Microsoft's Global Escalation Services team. He specializes in advanced Windows debugging and performance-related issues. For information about Windows debugging, visit blogs.msdn.com/ntdebugging.

Special thanks to Matthew Justice, a Microsoft software development engineer, who contributed to this article.

READER TO READER

Get Your IT Resume to the Top of the Pile

I was unhappy at my job and wanted to make a change. Even in these tough economic times, I knew that some businesses were still hiring. There were plenty of positions advertised on such websites as Monster and Dice, so I started submitting resumes for jobs that I thought were a good fit for my level of IT experience. However, I wasn't getting any responses.

I had a feeling deep down that my resumes were being overlooked, so I started looking for new ways to attract attention to them. I read many articles that gave tips on how to get a resume noticed, but they seemed to be just common procedural practices (e.g., compose a cover letter highlighting your experience in the field, use quality resume paper, emphasize certifications) that other candidates were probably also following, thereby making them the norm. I needed something to entice HR to read my resume and discover my skill set and abilities. With so many unemployed IT professionals submitting resumes, I knew that this was going to be a challenge.

I thought back to what I had accomplished in the past year at my place of employment, which was a manufacturing company in the automotive industry. For the past few years, the company had been experiencing tumultuous economic times. As part of my job, I often had to complete IT projects and solve system problems that would normally involve buying solutions. However, there was little money in the IT budget, so I was forced to think of creative, inexpensive ways to satisfy project requirements and solve system problems. I tallied all the money I had saved the company while still reaching department goals, without sacrificing quality of service. It was

a considerable amount—much more than my salary in fact.

That's when it hit me. In these slow economic times, HR managers are intimately familiar with their companies' need to reduce expenditures. And in these times, HR managers perusing large piles of resumes expect to read about candidates' great expectations and great skill sets. However, HR managers usually don't expect to read about a candidate successfully completing multiple projects on a shoestring in addition to having a great skill set. So, I revised my cover letter to reflect this way of thinking. First and foremost, I emphasized the amount of money I had saved the company by coming up with creative solutions. I cited several examples of the most beneficial projects that had the largest savings. Then, I revised my resume so it included information about these projects. I also tailored each resume to reflect the details of the position I was applying for.

Not wanting to have to compete with thousands of applicants like you do when you apply for jobs advertised on large websites such as Monster and Dice, I decided to find a better place to search for open positions to increase my odds. I also wanted a website that included local businesses, so I opted to look on Craigslist. The number of positions was significantly lower, but so were the number of applicants. I was now the big fish in a small pond.

After I changed my cover letter and resume, I started getting calls. I had sent out only five resumes with this new style and within 2 weeks I had received several calls and two interviews. One of those interviews led to a job offer, which I accepted.

—Matthew Kocot, IT manager,
Walton Engineering

InstantDoc ID 101756

Tell the IT community about the free tools you use, your solutions to problems, or the discoveries you've made. Email your contributions to r2r@windowsitpro.com.

If we print your submission, you'll get \$100.

Submissions and listings are available online at www.windowsitpro.com.
Enter the InstantDoc ID in the InstantDoc ID text box.

TOOL TIME

windowsitpro.com

Encrypt Files with AxCrypt

There are many free tools to encrypt sensitive data. Most of them let you encrypt an entire drive (e.g., Windows Vista's built-in feature) or create and encrypt a virtual drive (e.g., TrueCrypt). However, when you copy a file to removable media, the file isn't encrypted any more. To avoid this situation, I use Axantum Software's **AxCrypt** (www.axantum.com). You can use this free utility to encrypt a single file or a group of files. It remembers the file extension, so you can open the file later with a double-click.

AxCrypt integrates with Windows Explorer, so the easiest way to encrypt a file is to right-click it and select Encrypt from the AxCrypt context menu. When you encrypt your first file, you need to enter and verify a passphrase. If you select the *Use as default for encryption* check box, you won't have to enter a passphrase when encrypting other files during that session. After you click OK, the file is compressed, encrypted, and renamed using the format *Filename-Extension.axx*, where *Filename* is the file's original name and *Extension* is its original extension. For example, a file named *My review.doc* will be encrypted and renamed *My review.doc.axx*.

The AxCrypt context menu offers other options, including:

- **Encrypt a copy:** Makes a copy of a file, then compresses and encrypts the copy.
- **Encrypt copy to .EXE:** Creates self-decrypting .exe files that you can send to users who don't have AxCrypt installed.
- **Shred and Delete:** Overwrites a file's contents with random data, then deletes it. You can even use this option on files that aren't encrypted.
- **Clear Passphrase Memory:** Immediately clears the passphrases from memory, without waiting for logoff or reboot.

AxCrypt works on Windows 2000 and later. It can be used with files stored in a TrueCrypt volume for double protection.

—Serge Bedard, technology architecture specialist, CSST Quebec

InstantDoc ID 101651

■ Virtualization
■ Outlook

■ Domains

ANSWERS TO YOUR QUESTIONS

Q: I'm virtualizing my environment onto a small number of Hyper-V servers that are clustered to offer a highly available service. Where should I place my virtual domain controllers (DCs)?

A: A Windows failover cluster relies on Active Directory (AD) being available to offer services. You need to make sure that you don't place the virtual DCs in such a way that the virtual machines (VMs) can't start without the cluster being available, which in turn can't start without AD being available.

My advice is to place the configuration and virtual hard disks for at least two DCs on either local storage of each node or, if on a SAN, on storage that isn't cluster storage. Your DCs should be on at least two separate physical servers, so place one virtual DC on each of two Hyper-V servers. Don't place the DC resources on Cluster Shared Volumes (CSVs), because CSVs aren't available without the cluster, which isn't available without AD. Don't make the DCs cluster resources.

The DCs should be local VMs and you should always have at least two DCs in any environment for redundancy, in case one DC becomes unavailable or corrupt. You can then virtualize the other servers in your environment on CSV storage, but you've

ensured that if a single node fails or if the cluster can't make quorum, at least one DC is always available as a local resource.

Another option is to have at least one additional DC on a physical box or as a VM on another Hyper-V server that isn't part of the cluster. Just make sure you don't place all your DCs in one basket when that basket is part of a single failover cluster. You can obviously add additional DCs, and these could be on cluster storage.

—John Savill
InstantDoc ID 101630

Q: How can I view the complete header of POP3 messages?

A: Inbound messages from the Internet to Microsoft Office Outlook with POP have associated Internet headers outlining the source and routing information about that message. To troubleshoot the validity of a message or server-side antispy assessment, start by checking the message headers. The headers shown in Outlook don't include the message content by default. Starting with Outlook 2003, a registry change can be applied to provide more complete information about an email message.

In the Registry Key HKEY_CURRENT_USER\Software\Microsoft\Office\[version]\Outlook\Options\Mail, add the DWORD "SaveAllMIMENotJustHeaders" and assign a value of 1. The [version] is 11.0 for 2003 and 12.0 for 2007.

Outlook reads this key on start up, so to make the change you'll need to restart Outlook. In Outlook 2007, you'll find header content by expanding the Options ribbon object on a message. This opens the Message Options dialog box shown in Figure 1. For remote clients using POP3, the adminis-

Q: Where is Netdom in Windows 7?

A: Netdom is used in versions of Windows before Windows 7 for command-line domain membership tasks such as joining a computer to a domain. Windows 7 instead provides the PowerShell cmdlet Add-Computer, which allows you to add a computer to a domain or workgroup. The syntax for the command is

Add-Computer <domain name>

The cmdlet allows far more complex execution than Netdom. You can use items such as the organizational unit location of the computer account, credentials, and computer name in the command. Run the command Get-Help Add-Computer for all the syntax options.

—John Savill
InstantDoc ID 101628

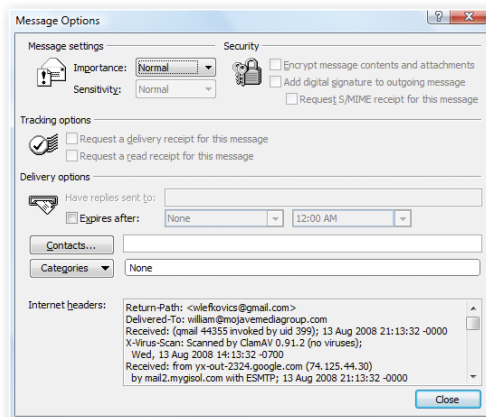


Figure 1: The Message Options dialog box

trator can request the entire message source to assist in troubleshooting measures.

This setting puts the entire message source in the same place where previously only the message headers were shown. If you make this change, messages stored in .pst files will take up more room. The header content now contains another representation of the message, which needs to be stored. It applies only to new inbound messages and does not add the message source to existing messages already downloaded from the server.

—William Lefkovich
InstantDoc ID 101057



William Lefkovich | william@mojavemediagroup.com
John Savill | jsavill@windowsitpro.com

[your one-stop resource for conference activities, products and services]

tech·ed | spotlight
north america 2009

Join Us at Tech·Ed 2009

TechEd is one of the most significant IT conferences of the year, and TechEd 2009 is no exception. *Windows IT Pro* and *SQL Server Magazine* will have a significant presence at the show. Here are some highlights to watch for:

Win a Prize at the *Windows IT Pro* Booth!

Be sure to stop by the *Windows IT Pro/SQL Server Magazine* booth (#411) to enter a contest and to chat with some of our editors and authors. We always like to hear feedback from readers, so let us know what you like (and don't like) about our coverage.

Author Sessions and Roundtables



A few of our authors are giving presentations at the show, including Senior Contributing Editor Mark Minasi (with sessions on Windows Server 2008 R2 AD features, Windows Kerberos, and Security with UAC/WIL) and Rhonda Layfield (presenting a Windows 7 from A to Z preconference session with Mark Minasi). *SQL Server Magazine* Contributing Editor Kalen Delaney will deliver a session about solving real-world DBA issues.

2009 Best of TechEd Attendees' Pick Awards

You'll also want to cast your vote in the 2009 Best of TechEd Attendees' Pick Awards, which lets TechEd attendees pick their favorite products on display at the show. You'll find voting kiosks scattered throughout the show floor that you can use to log on to the contest website and cast your ballot.



Live Blogs and Twitter Feeds

We'll be covering the show with some live blogs and Twitter feeds, so be sure to bookmark the *WindowsITPro.com* and *SQLMag.com* websites and follow our Twitter accounts for updates:

- **Best of TechEd Awards:** www.twitter.com/bestoftech09
- ***Windows IT Pro*:** www.twitter.com/Windowsitpro
- ***SQL Server Magazine*:** www.twitter.com/SQLServerMag
- **Jeff James:** www.twitter.com/jeffjames3
- **Amy Eisenberg:** www.twitter.com/witproamy
- **Sheila Molnar:** www.twitter.com/sqlmagsheila



Unleashing the power of SharePoint™

Tech-Ed North America 2009
May 11-15 | Los Angeles, CA

Win a Ducati



Visit us at

Booth 555



Game on - the pursuit for gold begins ...



Specops® Software™

Join us at Tech·Ed

Visit our Tech·Ed booth # 211 for a chance to win real GOLD!

What's New in DNS and Name Resolution?

New DNS features in Windows Server 2008 and Windows Vista—and upgrades to Windows XP and Windows Server 2003—get you closer to a world without WINS

by Mark Minasi

Considering the past two major Windows Server releases, we've come to expect changes to Windows name resolution and its primary naming protocol, DNS. Windows Server 2008 is no exception. In the interest of letting more people finally achieve an environment free of bandwidth-hogging WINS, the latest name-resolution upgrades permit a local-only version of DNS, provide a DNS zone that somewhat replaces the need for WINS, and boast two improvements to the way Windows systems find domain controllers (DCs). But before we get into the particulars, you need to ask yourself ...

Do I Need WINS?

The first question in any discussion about Windows name resolution is, "Can I disable NetBIOS and WINS now?" The answer is pretty much the same as it has been for the past nine years: "It depends." Essentially, you're stuck with WINS if your OSs or your applications need NetBIOS. If your OSs are Windows XP or Windows Server 2003, you're mostly fine from the OS point of view (although *ad hoc* networks might still need NetBIOS, as you'll see). But many of our organizations' applications depend on NetBIOS/WINS.

Unlike many aspects of Windows networking, here's a case where many small-to-midsized businesses (SMBs) usually have an easier time cutting the WINS cord than large enterprises do. Many small shops have just one site, run only Windows Vista or XP on their desktops, run only Windows Server 2008 or Windows 2003 in the back office, and use little in the way of applications beyond Acrobat Reader, Internet Explorer (IE), and Office—a

perfect environment in which to disable WINS and NetBIOS over TCP/IP. By contrast, virtually every large organization that I've encountered in my consulting experience depends on a collection of old, home-grown apps—apps that need WINS and will always need it, unless someone rewrites those apps (or unless an interesting new WINS workaround called GlobalNames can do the job—more on that later).

So, can you pull the WINS plug? The only way to be 100 percent sure is to thoroughly test your apps in a Server 2008-based network. That's a lot of work for most people. One way to help determine whether you still need WINS to run Performance Monitor on your WINS server comes from Microsoft coder Tim Rains in a 2004 blog posting called "Why you still run Windows Internet Naming Service (WINS)" (blogs.msdn.com/tim_rains/archive/2004/10/05/238236.aspx), who suggests logging the frequency of name-resolution requests to the WINS service. Perhaps the easiest answer, then, is to peek monthly at the Performance Monitor results, and when they start to dip, start looking more seriously into turning off WINS.

Replacing NetBIOS Broadcasts

A world without NetBIOS poses an important question: How can systems on the same subnet resolve names if they're not in DNS? Small *ad hoc* networks, test networks, and class lab networks sometimes either don't (or *can't*) include a DNS server—for example, most home networks don't include even a single copy of Windows Server, and Microsoft doesn't offer a DNS server for XP. Currently, any DNS-less network handles name resolution through NetBIOS broadcasts—bandwidth-wasting chatter that can chew up 10 to 20 percent of network capacity. Networks consisting of Server 2008 and Vista systems, however, can resolve names locally without broadcasts by using Link-Local Multicast Name Resolution (LLMNR), explained in RFC 4795 (tools.ietf.org/html/rfc4795).

LLMNR defines a multicast group at address 224.0.0.252 that LLMNR-capable systems (i.e., Vista and later—Microsoft hasn't offered a hotfix to add LLMNR to XP and 2003) can use to query names and get IP addresses in return. Thus, a query for the IP address of a system named PC33 doesn't

hammer every IP-using box on the network but rather those systems that have joined the multicast group. As you'd guess, you don't need the Computer Browser service to make this sort of name resolution work.

LLMNR's benefits don't stop there. Unlike NetBIOS name-resolution broadcasts, LLMNR queries can return not just IPv4 addresses but also any IPv6 addresses that a given system owns. And if you think you'll have to learn some new obscure network protocol to troubleshoot an LLMNR problem, fear not: LLMNR queries are just familiar, standard DNS queries. And speaking of standards, if you work with Apple or Sun equipment, you might already be familiar with a similar protocol on those platforms to accomplish local name resolution without a DNS server: multicast DNS (mDNS). LLMNR works very much like mDNS (and in fact I can't quite figure out why Microsoft didn't just use seven-year-old mDNS in the first place).

Network Discovery Rising

We have tools such as DNS, LLMNR, and NetBIOS not just for simple name resolution. In many cases, we're resolving a server's name to determine what services that server offers: what the names of its file and print shares are, what web services reside on it, and so on. Thus, name resolution is closely associated with the larger (and more valuable) topic of resource discovery.

In the late 1980s, Microsoft gave users on LAN Manager networks the ability to obtain a list of servers on their network and to query those servers for a list of their shares. That was Microsoft Resource Discovery 1.0—so to speak—and it was terribly primitive. (It required you to guess which server had the share you wanted, then look at the server's resources to see if you guessed correctly.) Later, Windows 2000 introduced the idea of publishing print shares and file shares in Active Directory (AD), with full keyword tagging and search capabilities. Unfortunately, I know of few organizations that benefit from resource publishing; for most networks, resource discovery boils down to asking someone else on your team, "Where are those files we use for the project, again?"

With Server 2008 and Vista—and XP, to a certain extent—Microsoft is trying yet another approach to resource discovery

Learning Path

WINDOWS IT PRO RESOURCES:

- "A DNS Primer," InstantDoc ID 7733
- "Deconstructing DNS," InstantDoc ID 48527
- "DNS Annoyances," InstantDoc ID 94456
- "DNS Configuration Errors Breed AD Horror," InstantDoc ID 43582
- "Solving DNS Problems," InstantDoc ID 39771
- "Split-Brain DNS," InstantDoc ID 99772
- "Troubleshooting DNS-Related AD Logon Problems, Part 1," InstantDoc ID 22774
- "Troubleshooting DNS-Related AD Logon Problems, Part 2," InstantDoc ID 23565

MICROSOFT RESOURCES:

- "Exchange Server 2003 and Exchange 2000 Server require NetBIOS name resolution for full functionality"
<http://support.microsoft.com/kb/837391>
- "Unique NetBIOS names must be used with WINS in an Active Directory forest with Windows 2000 Server and Windows Server 2003"
<http://support.microsoft.com/kb/927070>

with something called Network Discovery, which isn't a single protocol but rather a collection of protocols. Under Network Discovery, your system can use either the NetBIOS-based Computer Browser system or an LLMNR-based multicast system to enumerate other local systems. But to discover resources (e.g., who's sharing a color printer, who has the share related to the annual report), Server 2008, Vista, and XP—when equipped with SP3 and the hotfix from the Microsoft article "Network Map in Windows Vista does not display computers that are running Windows XP" (support.microsoft.com/kb/922120)—can exploit Web Services Discovery (WSD). Greatly simplified, WSD is the next generation of Simplified Services Discovery Protocol (SSDP) and Universal Plug and Play (UPnP), but with added security.

WSD is another peer-to-peer way for systems in a network to advertise their capabilities. The idea is that any system needing, for example, a shared printer can search for one by querying other systems on the network via a multicast message. Similarly, the system can send a directed message to another system, querying whether that system is sharing a printer. No, it's not a client/server approach such as AD publishing, but it's intended to be more standards-

May 2009

The Essential Guide to **Desktop** Virtualization

By Susan Perschke

UNISYS

As enterprise information networks encompass more line-of-business applications, the desktop computers used to access those applications become increasingly critical resources. Each desktop contains an environment customized for one user and that user's applications. The desktop environment often includes files containing sensitive information, as well as application and operating system software that must be regularly updated with security patches to ensure data integrity. And the desktop has information conduits—in the form of removable media and connection ports—that complicate keeping the desktop environment secure. The labor to maintain and secure desktops has increased dramatically over recent years, vastly increasing the cost of keeping enterprise desktops productive.

Many companies are turning to virtual desktop infrastructure (VDI), which moves data, processing and applications from the desktop onto shared, centrally managed hardware resources. In this Essential Guide we'll provide a technical overview of VDI and show how IT can leverage VDI to gain better flexibility and control while reducing TCO by as much as 30 percent.

Desktop virtualization concepts

Virtualization technology has been gaining popularity as a method to control server sprawl and, more recently, to streamline desktop provisioning and management. The traditional definition of a virtual machine (VM) is an efficient, isolated simulation of a real machine. One or more VM guests run within a virtualization environment on physical hardware termed the host. Because a VM simulates hardware it theoretically can execute any program written for hardware being simulated, such as a stand-alone application or an entire operating system (OS). Programs running inside a VM need not be compatible with the host hardware, although in modern VM systems this is usually a requirement to achieve optimal performance. In desktop virtualization the user's entire desktop environment, including applications, preferences, and storage runs on a server in the data center. The user accesses this desktop environment through a thin client terminal which remotes the keyboard, video, and mouse experience over the communications network. Together, the combination of the server executed VM and the thin client terminal make up a VDI.

In real-world deployments, VDI consists of multiple users operating multiple thin clients connected to multiple VMs hosted on several data center-based servers. VDI allows IT departments to provision and manage end-user applications and desktops from groups of servers in the data center, thus reducing the number of desk side support visits and providing better control over desktop management tasks such as deploying software updates and patches. The enterprise also gains tighter control over security by limiting access to removable media and other sources of data leakage; the thin client can simulate CDs and DVDs when needed, and its network connection is restricted to reach only the VDI servers. VDI's thin client architecture lets IT "lock down" access points and prevent the installation of unsupported software. Yet VDI gives users the flexibility of "anywhere, anytime" access to data and applications, all within their familiar desktop environment.

Understanding Thin Client Architecture

Thin clients do not process data in the traditional sense, but instead rely on a server or group of servers (sometimes referred to as a server farm) for processing and data storage. The thin client provides a graphical interface, through its physical keyboard, mouse, and screen, to desktop applications residing on one or more virtual servers (Figure 1). Thin clients may be as simple as a "smart" keyboard-video-mouse Ethernet appliance, or as complex as thin client agent or terminal emulator software that simulates a thin client appliance on legacy PC hardware. Either approach lets users access their target application or unique desktop environment hosted virtually on local or remote VDI servers.




Figure 1 - Virtualized Desktop Solution Architecture

Virtualization and management servers integrate into the existing data center environment and communicate with thin client devices at end users' workstations to create a seamless, manageable, and secure virtualized environment.

There need not be a one-to-one correspondence between thin clients and VMs. An individual thin client can be used to access multiple VMs, either because the thin client is used by multiple users at different times or because an individual user may log into different VMs to accomplish different tasks or to access different company networks. Likewise, an individual VM can be accessed by one or more thin clients at any point in time. Multiple users can execute on the same VM in the same way that multiple users can sign into a single PC through network connections. Each VM can be provisioned with the amount of processing power and memory that users of that VM typically require. This avoids the significant over-provisioning that can otherwise occur when all users are given physical PCs that meet the needs of the organization's most demanding power users.

Data Storage

With VDI data is usually stored centrally on dedicated storage subsystems such as RAID Direct Attached Storage (DAS), or RAID-based fibre channel or iSCSI Storage Area Network (SAN) devices. This approach increases both the security and availability of the data because a centralized, redundant data storage infrastructure is more fault-tolerant, easier to secure and back up, and generally better able to support the entire organization. Such storage is sometimes termed "virtualized" because aggregated storage servers appear as multiple simulated DAS devices to host and virtual machines.

What are the Benefits of Desktop Virtualization?

Simplified Management

Most IT departments spend substantial resources deploying and maintaining desktop computers. Challenges include making sure each desktop computer has the cor-

Unisys Desktop Virtualization Solution Can Help Reduce TCO

Q How will the Unisys desktop virtualization solution help reduce my Total Cost of Ownership (TCO)?

A Unisys Consolidated Desktop Solution (CDS) addresses the growing—and expensive—challenges associated with managing and protecting Windows desktop environments while significantly decreasing the associated costs. By virtualizing and hosting the processing and data on centralized servers and storage, CDS eliminates the complexities and risks synonymous with decentralized computing at the desktop level. This brings a greater level of control to the desktop model, which is typically inadequately secured, expensive to procure and maintain, and extremely complex to manage. As such, CDS helps organizations better control access to sensitive data, move green IT across the productivity chain, and address a myriad of regulatory compliance issues and administrative concerns – all while reducing the overall TCO of the Windows based desktop infrastructure by 30 percent or more.

Q Security is paramount in our organization; how does your solution address this critical issue?

A Implementing any VDI solution will definitely improve your desktop security by moving data off of your PCs and into the data center where it is protected behind a corporate firewall, cared for by centralized backup processes and guarded by a disaster recovery plan. Unisys pays special attention to the security aspect of VDI and offers a host pooling capability that safeguards against server failure. On the client side our model 2140 thin client can convert an existing PC into a secure terminal that provides added protection over and above simply running terminal emulation software on your PC.

Q We have limited space for additional hardware—are we going to be able to use some of our existing infrastructure with a Unisys solution?

A Certainly—there are several ways to leverage existing infrastructure with Unisys solutions. On the server side, you may be able to use server consolidation to reduce the number of servers you need and free up resources that can then be repurposed for your VDI environment. On the desktop side, we have several thin client alternatives that allow our customers to reuse existing devices. The Unisys Thin Client Model 4140 replaces your end user's existing PC chassis, but in most cases allows them to reuse their monitor as well as keyboard and mouse. This provides for some level of investment protection by allowing you to benefit from the value of virtualization without replacing your end users' entire desktop. We also offer a very cost-efficient client option, the Unisys Thin

Client Model 2140 secure stick. It allows your organization to fully leverage existing desktop investments into a virtualized environment. The 2140 is a USB stick that supplies a boot image that transforms existing PCs and laptops into a secure terminal.

Q We have spent substantial resources on training and certifying our staff on Microsoft products—are we going to be able to leverage this knowledge or will there be a substantial learning curve for our IT staff?

A The Unisys Consolidated Desktop Solution (CDS) is architected to leverage Microsoft technologies such as Windows Server 2008 with Hyper-V and the Microsoft System Center management tools. Our documented implementation of these and other Microsoft software technologies allows the use of existing skills within your organization, while speeding time to value. But our commitment to customer satisfaction doesn't end there. With Unisys as your partner, you have a one-stop-shop for assessing and configuring, designing, attaining, deploying, and maintaining an automated, dynamic, and adaptive Real-Time Infrastructure for the Windows desktop environment.

Q How can the Unisys virtual desktop solution help us meet our green IT initiatives?

A Whether your green IT initiative seeks to reduce energy consumption, carbon dioxide (CO2) emissions, or financial costs of operating IT, a Unisys CDS deployment can help. The Unisys Consolidated Desktop Solution helps IT organizations streamline, consolidate and centralize operations to improve asset utilization and realize cost and energy savings across the board. Our prescriptive architecture for deploying a VDI infrastructure provides cost savings by centralizing and streamlining IT operations, while at the same time allowing IT to reduce its carbon footprint and consume less power.



John Keller, Unisys Director of Desktop Infrastructure Solutions

John Keller is actively engaged as the Unisys Director of Desktop Infrastructure Solutions. His leadership role is defining, developing, and marketing the Unisys Consolidated Desktop Solution. Additionally, John served as both Marketing Manager and Product Manager for the Unisys Enterprise Server ES7000 program. Prior to his role in the US, John was the Program Manager for Enterprise NT for Unisys Switzerland. John initially joined Unisys Switzerland in 1991 as a Consultant for PCs, servers, and solutions. Former to this, John managed IT systems for several medium-sized enterprises in Switzerland.

rect mix of applications, plus the latest software patches, virus definitions and critical updates to keep users productive. It's not easy to ensure that any given change fully penetrates the installed desktop population. With the centralized approach of VDI, IT departments can ensure 100 percent penetration when deploying application software, patches, and updates. Building and deploying new virtual desktop guests can be accomplished in a matter of minutes instead of hours or days.

Easier Support and Maintenance

Many desktop support calls are spent solving problems that don't necessarily have much to do with the end user's job, but rather are brought on by installation of personal software or malware introduced by recreational Web browsing. Furthermore, because traditional PCs can be quite prone to hardware issues, support personnel may find themselves troubleshooting and replacing faulty or failing hardware.

With thin clients, you have fewer moving parts, lower incidence of hardware problems and consequently, fewer desk side support visits. VDI's centralized administration also means faster, more reliable deployment of only pre-tested and certified software components. The IT department gains better control over its IT assets and can provide better support to end users.

Improved Availability and Security

A centralized structure gives you much better control over your IT organization, its assets, and applications. Thin clients let you lock down or remove thin client access points as needed. Securing access points helps prevent unwanted content such as viruses and unsupported applications from being introduced into the IT infrastructure, and also provides a single point of control for limiting access when job assignments change.

Centralized management of data also ensures better compliance with corporate policies and regulatory requirements such as HIPAA and SOX. You can guarantee backup integrity because the backup process isn't dependant on the user leaving their machine turned on. You can enforce security compliance because the thin client can be constrained to a separate, isolated network that doesn't have generalized access to enterprise compute assets, such as databases and Internet connections.

Another major benefit of a virtual desktop is that users can access their desktop from anywhere they have access to a thin client or terminal emulation software. Since the desktop and data are now hosted centrally, a telecommuting associate can access his desktop from his home PC or a traveling sales person can access the home office from a laptop or hotel computer without compromising data stored safely on the server—data which never really leaves the four walls of the corporation.

Finally, VDI can potentially simplify disaster recovery and business continuity processes. If your backup system can replicate VDI servers on virtualization hosts in a DR data center, you can theoretically move your entire user population to a new location and make them immediately productive with basic thin client hardware or software.

Improved End-User Productivity

In addition to improving the productivity of the IT staff with fewer desktop support visits, VDI can greatly improve the productivity of end users. End-user downtime due to desktop configuration problems, software problems, or

hardware failure can be costly to the organization, particularly when data loss occurs. Recovery time for a physical desktop hardware failure can be hours to days. VDI with redundant failover can help eliminate computer failures, and end-user equipment failures require just minutes to repair, thus boosting end-user productivity.

In a virtualized desktop environment users can access server-based applications using the same desktop environment they are familiar with on the local area network. Virtualization also lets you seamlessly add disk, memory, and CPU capacity when needed, and shift workloads dynamically to accommodate changing demand. Compared with the old hardware upgrade process, which could take months to accomplish, VDI expansion is instantaneous.

Meeting Environmental Mandates

Going green and shrinking the carbon footprint are concepts that are taking hold in IT as well as in organizations as a whole. With increasing energy costs, the cost of powering and cooling a data center that has multiple servers in addition to hundreds, if not thousands, of desktop computers can add up in a hurry. By using energy-efficient thin clients and better utilization of server hardware, organizations can save a bundle on energy costs. The concentration of end user computing hardware lets you leverage economies of scale that drastically reduce energy consumption. A typical thin client appliance, for instance, consumes 25 watts including the LCD screen, one tenth the 250 watts guzzled by a traditional PC. Because a VDI server hosts many users simultaneously, it makes much more efficient use of data center computing horsepower, consuming as little as 50 watts per active user and nothing at all for inactive ones. A 1000-user enterprise could thus cut its desktop power and cooling costs by 70 percent.

Business and IT Benefits

VDI improves the quality of IT services and provides more predictable operational availability and costs. Streamlining the desktop environment through virtualization frees up IT personnel to focus on higher value tasks within the organization. Power and equipment savings feed money back into strained IT budgets, and makes VDI nearly self-funding.

As corporate governance becomes more demanding, VDI's centralization simplifies compliance by keeping all sensitive data in one place under tighter control. VDI also simplifies auditing tasks required to prove compliance, since the audit process need not involve end user cooperation. And by restricting all application communications to a single encrypted tunnel per client, you achieve a single point of control for data security, reducing vulnerabilities and the chance that a minor administrative error could enable a serious security breach.

By taking advantage of the benefits of VDI, business becomes more agile and responsive to changing technology and business needs. Through consolidation and efficiency improvements, organizations can realize TCO savings of 30 percent or more.

Choosing the right solution

Selecting a systems integration vendor

It is important to select a systems integration vendor that offers the right solution for your organization. Look for a vendor with experience in the areas of desktop management, consolidation, and virtualization that can provide a solution customized to your environment. The vendor also

Top 10

Reasons to Implement Desktop Virtualization

10 Accommodating Access Needs – Managing remote access for traveling and remote workers poses considerable management and access control challenges. With desktop virtualization users can securely connect to their desktop from remote locations using a variety of client devices. This provides a uniform and familiar desktop environment for remote users regardless of their day-to-day location. Data security is vastly improved because it is stored in the data center; remote users don't have to transport sensitive enterprise data on their laptops or memory sticks. Desktop Virtualization also enhances the ability to provide secure non-employee access to select corporate data and applications without giving contract workers free access to the entire enterprise.

9 Better Resource Utilization – Virtualization provides better resource utilization from both a hardware and software perspective. You can extend existing desktop computer life through virtualization because the processing is moved from the local desktop to the server. You can better utilize software licenses because you only need licenses for the actual number of desktops used. Green initiatives may be better served by virtual desktops because you can host multiple virtual desktop environments on a single, higher-availability server. This approach can also reduce energy consumption.

8 Richer Application Delivery – Delivering an efficient and rich desktop experience to end users requires time and resources that may be in short supply in IT departments operating on more limited budgets. By virtualizing the desktop environment, IT departments can provide a customized desktop environment for each end user from a centralized management point. End users have quicker access to new and improved technology without IT making a service call to each user's cubicle.

7 Meeting Environmental Mandates – Going green and shrinking the carbon footprint are concepts that are taking hold in IT as well as in organizations as a whole. Computers, from laptops to large servers, generate a lot of heat and use substantial amounts of electricity. By moving to desktop virtualization, IT organizations can eliminate the need for energy-demanding desktop computers and move to more energy-efficient thin clients. Processing tasks are moved to the data center where one or several virtualization servers can host hundreds of virtual machines.

6 Improved Regulatory Compliance – Desktop virtualization helps IT organizations ensure that data guidelines are in line with regulatory mandates such as HIPAA and Sarbanes-Oxley.

5 Improved end-user productivity – User downtime due to desktop configuration problems, software problems, or hardware failure can be costly to the organization. Desktop virtualization with redundant failover can help eliminate computer failures, thus boosting end-user productivity. Furthermore upgrades to new operating systems can be performed at the server without needing to replace older desktop hardware or take the user offline for days.

4 Easier Support and Maintenance – As end users require a more sophisticated and personalized desktop environment they become more challenging to support and maintain. Most IT departments spend substantial resources deploying and maintaining desktop computers. Challenges include making sure each desktop computer has the latest software patches, virus definitions, and critical updates to keep users productive. With the centralized approach of virtualization IT departments can ensure 100 percent saturation when deploying software patches and updates.

3 Simplified Management – Managing desktop computers tends to be one of the costliest and time-consuming responsibilities of IT management. A virtualized desktop environment can greatly simplify desktop management tasks by centralizing the client operating system and application management. With virtualization a new client can be deployed in minutes instead of the hours it usually takes to configure a desktop computer from scratch. Virtual desktops are also easier to update, patch, and back up.

2 Improved security – Desktop virtualization allows IT organizations to improve control over data backup and security. User data will be centrally stored and backups can be performed without input from end users. Gone are the days when a user saves important data to their local hard drive and exposes the organization to potential loss either through hardware failure, loss, or theft. IT departments can ensure that software patches, critical updates, and virus definitions are deployed on schedule.

1 Reduced Cost – Lowering total cost of ownership (TCO) is the goal of every IT department. Virtualizing the desktop environment allows organizations to eliminate costly desktop computers and lease agreements by using thin clients that allow users to connect to their desktops hosted as virtual machines on a virtualization server. Thin clients allow for longer equipment lifecycles. You don't need to upgrade end-user equipment to accommodate a new more resource-intensive application because most, if not all, of the processing is performed on the server.

should be able to easily scale beyond the initial implementation. An experienced vendor will be able to create an end-to-end, pre-tested, and integrated desktop solution that

- Employs a simplified architecture and easily integrates into your existing infrastructure
- Provides scalable thin-client options that offer flexibility while preserving the end-user PC experience
- Uses state-of-the-art virtualization technology to replicate applications and processing in a secure state away from the end-user layer

Keep in mind that selecting a dependable vendor is as important as selecting the technology. You want to work with a vendor that has a proven track record and one that will respond quickly to address any problems that may arise. While many industry players offer one or more independent VDI components, you may need a lot of in-house expertise to implement their solutions. Alternatively, you may want to consider a prescriptive VDI solution delivered with pre-built, pre-certified components designed to work together from the outset. Vendors such as Unisys offer an end-to-end VDI solution that allows you to quickly achieve the centralization and control needed to move up the ladder to a more dynamic environment.

Cost and Licensing Models

When selecting a VDI solution, look for vendors and products that fit in with your virtualization goals and existing IT infrastructure, not only from a hardware and software point of view, but also from a budgeting and licensing standpoint. Many vendors provide calculators that will allow you to calculate Return on Investment (ROI) and TCO. As with any major purchase evaluation you want to make sure you take licensing costs into consideration with other capital expenditures to get an accurate cost and savings picture.

Supported Platforms and Management tools

Be sure to select a solution that supports your IT software management platform, because this will enable you to leverage existing applications using existing technician skill sets. Management tools should provide centralized storage, backup, and image maintenance, as well as performance monitoring and event notification. Although each of the major VDI solution manufacturers Citrix, Microsoft, and VMware offer their own native management solutions, you may want to consider third-party management products that are platform-agnostic,

Deploying a Desktop Virtualization Solution Planning and design

One of the challenges of VDI is to properly plan a deployment so you provision the right amount of infrastructure to support the organization while deriving the maximum benefits from virtualization technology. Capacity planning, a key element of virtualization planning, involves determining how many workloads you can run on one physical server without degrading performance or introducing operational risks. When planning for a virtualization deployment start by taking inventory of the infrastructure you already have in the data center. Many organizations have tools such as Microsoft System Center in place to map out IT infrastructure.

With this information in hand it's time to look at server utilization. Servers tend to be underutilized in a non-virtualized environment, but by monitoring your IT infrastructure over a period of time under typical load conditions, you can deter-

mine basic utilization levels. Servers with low current utilization levels may be good candidates for server consolidation, thus freeing up systems to become desktop virtualization hosts. Consider how you can leverage current infrastructure together with the capital investment that may be needed to fully implement a thin client strategy.

Additional considerations when planning for virtualization:

- Are you planning to upgrade to a new operating system (Windows XP to Vista, Windows Server 2003 to Server 2008)?
- Will you still need to support non-virtual desktops?
- Do you have an adequate centralized storage platform to migrate data from desktop hard drives?
- Have you planned for adequate LAN (and WAN, if necessary) bandwidth to minimize latency between thin clients and upstream application servers?
- How will you handle disaster recovery?
- Will your data center's available power, cooling, and rack space be sufficient?

On the desktop side, be sure to inventory the various applications currently in use and have a plan for migrating and load balancing these applications in a server-centric environment. Desktop applications may behave differently on virtualized server hosts, so have a comprehensive testing cycle in your plan.

Implementation

In the deployment phase, starting your VDI project with a few pilot users is always a good strategy. The pilot group provides a test environment that can be very helpful in rooting out system problems and bottlenecks. Once things are running smoothly with the pilot group, you can implement VDI across a broader scope of users within the organization. Rolling out large deployments in phases also helps manage user expectations while keeping the organization running smoothly.

Post-implementation

As with any major project you may want to perform a post-implementation audit once the VDI project is completed. A post-implementation review will give you valuable information about your project such as the impact on the business (pros and cons), ROI, bottlenecks that need to be resolved, and areas where you and your IT staff might improve when embarking on future projects. You'll also be able to verify projected cost savings and use them to reduce the TCO for your VDI project.

Wrap Up

In this Essential Guide we've seen how thin client VDI can allow your IT organization to meet a number of goals such as simplified management, better utilization of IT infrastructure, reduced power consumption, enhanced resilience, increased data integrity, and tighter security all while providing improved ROI and reducing TCO. Virtualization can greatly improve the efficiency and availability of applications in the data center. With VDI, users are able to access a broad array of applications that you can tailor to meet individual needs. Business becomes more agile and responsive, allowing IT decision makers and planners to shift focus from rote hardware and software management to higher value business goals.

Susan Perschke has more than 20 years of experience as the lead programmer and executive manager of a database and web technology firm. Her experience includes consolidation through virtualization and remote desktop systems.

based than NetBIOS browsers because it's built atop the whole notion of web services, Simple Object Access Protocol (SOAP), and so on. You can recognize one kind of web services discovery traffic on your network when you see UDP messages sent to port 3702 on multicast address 239.255.255.250. Those messages—which are SOAP-formatted XML—are general queries seeking a gateway for connecting to the Internet. You'll also see TCP messages directed to particular servers on either port 5357 or port 5358 (one is HTTP, and the other is HTTPS, depending on how the client authenticates to the server); these messages are called *directed* discovery messages because the system isn't querying the multicast group but rather a particular server.

What sort of systems can participate in Network Discovery and WSD? Clearly Server 2008, Vista, and later versions of Windows systems can. To a limited extent, XP systems can participate—but only if you keep NetBIOS over TCP enabled on all systems, run SP3, and install the hotfix from the previously mentioned Microsoft article.

Client-Side DNS Changes

Some of Server 2008's most important DNS changes appear not in the DNS server software but in the DNS client software. In AD, member systems must first find a DC before they can log on to a domain. A member system uses DNS to do that, and once it

mand to force the system to find a new (and presumably closer) DC.

With Server 2008, Vista, XP, and Windows 2003 (the latter two patched with the hotfix), the system no longer remains with a given DC until rebooted. Instead, it remembers its current "DC buddy" for 12 hours, then re-queries DNS when it needs a DC after that point. You can modify that behavior to make the time value larger or smaller: Navigate to the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters registry subkey, create a REG_DWORD entry called ForceRediscoveryInterval, and set the interval value in seconds. It's a good feature, but again, you'll need the aforementioned hotfix to add it to XP and Windows 2003.

That's not the only good news in terms of ensuring that your members find nearby DCs. Server 2008 and Vista provide another change to DNS client-side DC-location behavior. When querying DNS to find the domain's DCs, AD clients from Win2K and later ask DNS for two lists of DCs. First, they ask DNS for the list of DCs in their site. If the member asks all local DCs to log it on but no DCs respond within 400 milliseconds, the member asks DNS for a second list—the list of all DCs in all sites—and the member requests each of those DCs to log the member on. That's the reason you sometimes get logged on by a DC in Outer Mongolia when you're sitting in Austin, Texas.

CurrentControlSet\Services\Netlogon\Parameters subkey, create a REG_DWORD value named TryNextClosestSite, set it to 1, and reboot. Or use the Group Policy setting at Computer Configuration\Administrative Templates\System\Netlogon\DC Locator DNS Records\Try Next Closest Site. Unfortunately, there's no retrofit available for XP or Windows 2003.

The GlobalNames Zone

Microsoft knows that many people are stuck with old, NetBIOS-centric applications that they can't replace anytime soon, and that many of those folks want desperately to move away from NetBIOS as soon as possible. For those people, Server 2008's DNS offers an answer: the GlobalNames zone.

The idea is that many NetBIOS-fixated systems will work just fine without WINS, as long as you feed those applications short server names: Tell one of these apps to communicate with server44.bigfirm.com, and it complains; configure it instead for just server44, and all's well. Again, this isn't the answer for all old applications, but it works for many of them, and those are the applications that are the target of GlobalNames.

As you might know, the reason the hostname works in DNS-dumb apps is that your client OS is configured with a DNS suffix such as bigfirm.com. When the DNS-dumb app asks your client OS to resolve server44, your client's DNS server automatically attaches that DNS suffix to the hostname and asks its DNS server to resolve a Fully Qualified Domain Name (FQDN) such as server44.bigfirm.com. DNS can handle a name like that, so it returns an IP address and the old application is happy.

That's one reason why single-domain forests are more likely to be able to turn off WINS and not need GlobalNames—the DNS client's action was just the right thing. But what about an enterprise that has two or more domains? Suppose our client had a DNS suffix of mmco.com and the server44 system had a DNS suffix of bigfirm.com. In that case, there's no guarantee that server44 happens to be in the same domain as the client that's trying to access it, so the DNS name resolution would fail. Yes, you can deploy a series of DNS suffixes via Group Policy, but apparently a significant number of Microsoft's customers found that hard to manage.

Some of Server 2008's most important DNS changes appear not in the DNS server software but in the DNS client software.

finds a DC, it continues to use only that DC for future authentications until the member is rebooted (if it's a pre-Vista AD member). Typically, that's sound functionality, but in some situations a member's initially discovered DC is distant, and a sluggish WAN link between a member and its DC can slow authentication—a problem that continues until a reboot occurs or until someone notices and uses the Netdom Reset com-

With Server 2008 and Vista, you can optionally change how your AD members look for a DC. If they've queried all local DCs and gotten no answer, you can tell your members to next query DNS not for the worldwide list of DCs but instead the DCs from the next nearest site (as computed by site link costs). This behavior isn't the default; to enable it, navigate to the HKEY_LOCAL_MACHINE\SYSTEM\

■ WHAT'S NEW IN DNS?

Server 2008-based DNS servers can, if specially configured, contain a special zone called GlobalNames, a single-label name unlike more common two-label names such as bigfirm.com. When asked to resolve server44, our hypothetical client would add mmco.com and thus query DNS for server44.mmco.com—a query that would fail. A Server 2008 DNS server configured with a GlobalNames zone, however, wouldn't give up yet. It would take the hostname (server44) and look in the GlobalNames zone in search of an answer to this single-label query. The salient aspect of the GlobalNames zone is that it can include only CNAME (alias) records—records you've statically created. If any of those CNAME records match server44, the DNS server looks at the target of the CNAME record—for example, server44.bigfirm.com—and queries its A record. That query should succeed; we've already posited that a server44.bigfirm.com exists, so the Server 2008 DNS server would return that IP

address to the client, and the old application would be happy.


Setting up GlobalNames is relatively easy, if stringent. First, all the DNS servers in your intranet need to be running Server 2008 because only a Server 2008 DNS server is smart enough to remove the old suffix and find the CNAME; if the client queries any other sort of DNS server, the game is over. Second, you'll probably find managing GlobalNames easiest if you make it an AD-integrated zone and put that zone in the ForestDNS partition. By default, Server 2008 DNS servers don't know how to use GlobalNames until you type the following command, which you'll need to do on each DNS server:

```
dnscmd /config /  
EnableGlobalNamesSupport 1
```

GlobalNames support in Server 2008 DNS has just one more (undocumented) quirk: It looks at the DNS suffix of the client requesting the name resolution and—if that DNS suffix doesn't match any zone that the

DNS server is authoritative for—the DNS server, in my experience, doesn't ever look in GlobalNames.

Worth a Look

If you're already rolling out Vista and Server 2008, give its new name-resolution features a look. By doing so, you'll understand how they've changed what goes on under the hood of your networks. GlobalNames might just make your life easier in regards to legacy apps. And if you're still running XP or 2003, don't overlook their optional name-resolution upgrades, either—they don't want to miss out on the lower-bandwidth fun! 

InstantDoc ID 101605



Mark Minasi

(www.minasi.com/gethelp) is a senior contributing editor for *Windows IT Pro*, an MCSE, and the author of 25 books, including *Administering Windows Vista Security: The Big Surprises* (Sybex). He writes and speaks around the world about Windows networking.

Even Meerkats Monitor.

Don't wait until it is too late,
start monitoring today.



AWARD-WINNING EVENT LOG MONITORING & CONSOLIDATION,
SYSTEM HEALTH, ENVIRONMENT AND NETWORK MONITORING SUITE.



© Copyright 2009 NETIKUS.NET Ltd. All Rights Reserved. EventSentry is a registered trademark of NETIKUS.NET Ltd in the United States and/or other countries. All other trademarks are the property of their respective owners.



“Customers will live in a hybrid world of some workloads living on-premises and some workloads living in the cloud. Our strategy is to provide protection for those workloads, whether they live on-premises or in the cloud.”

Doug Leland
Microsoft

Identity and Security: Microsoft's Next Generation

Protecting workloads on-premises and in the cloud by Jeff James

Identity, access, and security have always been top-of-mind topics for IT pros, but recent developments in hosted services, cloud computing, and Software as a Service (SaaS) have created challenges: How do you ensure the integrity of identity information in the cloud? How can you be sure that the right people are getting access to your vital corporate information in both on- and off-premises services?

Microsoft saw the writing on the wall in these areas and merged its Access and Security division with its Identity and Access division late last year, creating the Identity and Security Business Group. This merging of identity and security could mean Microsoft products and technologies such as Active Directory (AD), Windows Rights Management (WRM) Server, Active Directory Federation Services (ADFS), Microsoft Forefront, and Identity Lifecycle Manager (ILM) all might work more closely together in the future, making it easier for IT pros to deploy and manage their access and security infrastructures. To see what Microsoft has planned in this area, we recently spoke with Doug Leland, general manager for the Identity and Security Business Group.

Jeff James: What are your overall goals for the Identity and Security Business Group?

Doug Leland: Our overall goals are to provide identity and security solutions for the broadest range of customers out there, from some consumers all the way up to the largest enterprises, and provide a range of customer solutions from being able to protect their

■ IDENTITY AND SECURITY

endpoints—endpoint security—to being able to protect their strategic workloads—for example messaging and collaboration. At the same time we want to be able to provide unprecedented access to information applications and networks, all supported through a unified management experience across both identity and security.

Jeff James: What are some of the reasons why you think it's important to combine security and identity?

Doug Leland: I think the key drivers for us in bringing identity and security together are anchored in our customers' needs, and of course in the needs of our partners, who are ultimately providing those services to our customers. One of the things we've observed in talking to our customers and our partners is that the business needs around identity and security have been converging for years. We saw this convergence of business requirements and that dictated a need for us as a company to be able to solve these problems together.

Jeff James: Based on your market research and feedback from customers, what do you see as the top things IT pros are looking for help with in the security and identity areas?

Doug Leland: Compliance is certainly one of the key needs, and that's an area where we believe the identity and access solutions we provide help enormously. The second area is around business agility, which we think of as helping customers realize the benefits of business models or new ways of conducting business. The third area is around being able to do all this, to ensure compliance and ensure agility but to do it at the right cost, with effective cost benefit. Those are the key needs that we hear reflected again and again from our customer base.

Jeff James: Could you talk about Microsoft's current identity and security products and where you're heading in the future?

Doug Leland: In the identity and security space, there are a range of point solutions that are available in the

marketplace. And more and more as customers are investing in these point solutions, they are realizing that they're not really the best answer.

The problem with these solutions is primarily around cost—the cost of acquiring them, which tends to be at the higher end, and the cost of integrating them with the existing systems, and then ultimately the challenges associated with not having end-to-end visibility across those point solutions.

How can you be sure that the right people are getting access to your vital corporate information in both on- and off-premises services?

One of our strategies is to provide unification across identity and security management, so that through a single console an IT pro can both manage the implementation of identity and access management and also security management, and at the same time provide the end-to-end visibility that is needed to ensure the company is in compliance.

The second key aspect is delivering end-to-end access and end-to-end protection. At its core, security is all about keeping the bad guys out, and identity is all about letting the good guys in. That's why I call it the yin and the yang or two sides of the same coin.

Our strategy here is to deliver a set of solutions that provide that end-to-end access and protection. What we mean by end to end is that it's a multi-layered approach—from the network to the applications to the data—and ultimately providing both that protection and that identity-access layer in the stack, so to speak.

The first strategy is about extending the platform. We feel the best way to provide secure access to companies, and good end-to-end or secure end-to-end protection, is to be able to build these technologies into the core infrastructure, into the platform, that these companies are implementing.

But also to foster the development of a broad ecosystem of partners who are taking

advantage of these platform capabilities and delivering applications themselves that are inherently identity-aware and are more secure.

Jeff James: How does this product strategy work with things like OpenID, your own Sterling product, Cardspace, and other products?

Doug Leland: Interoperability and integration is a core piece of the strategy,

particularly when you think about an identity infrastructure, where identities need to be able to operate across a wide range of resources. Will those resources be within your organization?

It might be an application, website, or internal portal, but you might also have an employee or identity that needs access to resources outside your application, for collaborating with an organization or taking advantage of software delivered as a service (which, of course, Microsoft is now doing with our Business Productivity Online Services), where identity is critical to providing that foundation for authentication and access, secure authentication, and secure access of those services.


So interoperability becomes fundamental, and we've been working with the industry around a set of frameworks and a set of standards. OpenID is an example of a standard that we are working with, and it doesn't stop there. When you look at the platform capabilities that we're building around Active Directory, which supports LDAP, we're actively building in and supporting the core standards which allow for a high level of interoperability at the identity and security level.

Jeff James: Could you talk about Microsoft's relationship with RSA, working



Mediterranean Shipping Company has discovered a new form of energy.

Mediterranean Shipping Company (MSC) is the second-largest container ship line in the world, with a database that tracks more than 210 billion transactions a year. The company recently upgraded its database to Microsoft® SQL Server® 2008, not only to handle this massive load, but also to simplify MSC's database administration and help ensure high availability. Which is like a new form of energy for MSC. See the whole story at SQLServerEnergy.com



To get the full MSC story on your phone, snap a picture of this tag. (Requires a free mobile app from <http://gettag.mobi>)



Microsoft®
SQL Server® 2008

■ IDENTITY AND SECURITY

together to develop a modular approach to protecting information, and what you're doing with RSA?

Doug Leland: One of the key dialogues or challenges that customers are facing right now is protecting the information assets that they have as an organization, whether that be HBI (high business impact) data or PII (personally identifiable information) that we hold about many of our employees and/or the customers and businesses that we deal with.

And as we've seen with the rising publicity around data breaches over the last couple of months and even years, this problem is only growing and it's being exacerbated by the downsizing that's taking place.

Now you have the rise of a disgruntled employee who has easy access to the crown jewels of the organization, which is the information.

Given this backdrop, we saw the opportunity to again converge a set of needs around securing information, which has been approached via a market approach which is called data leakage protection.

And converge that approach with the enterprise ID management approach, which is all about providing identity-based access information, enabling customers to access information but access it securely and have those access privileges be part of the information itself.

So we reached across the aisle to one of our key partners, EMC or RSA (the security division of EMC), to partner at a technology level and a sales and marketing perspective to deliver a unified solution across the classic DLP and the enterprise rights management space, to build a more comprehensive solution that addresses these broader-range needs for securing the information and providing access to the information.

Jeff James: We've heard from readers concerned not only about security and identity in the cloud but also between the cloud and their own on-premises environments. How do you address IT pros' concerns?

Doug Leland: We're hearing the same thing from customers, in terms of their desire to take advantage of the cost benefits and economics of being able to operate in a Software+Services environ-

ment where they have a choice of running workloads either on-premises, or in the cloud, or some combination of both. And we believe from the company perspective that it's an "and" versus "or."

In other words, we will deliver solutions for use on premises, but also in the cloud, and those need to be able to easily migrate back and forth, and also to interoperate, meaning customers will live in a hybrid world of some workloads living on premises and some workloads living in the cloud. Our strategy is to provide protection for those workloads, whether they live on-premises or in the cloud.

A couple of examples: Today, when a customer purchases the Business Productivity Online Suite from Microsoft, it comes protected by Forefront. So, specifically when a customer buys SharePoint Online or Exchange Online, those come already protected with their companion Forefront products, Forefront Security for Exchange, or Forefront Security for SharePoint.

That is a model we will continue to follow, and we will also build out what you may think of as standalone offerings for cloud-based protection of either non-BPOS workloads or protection of on-premises solutions. A key example already available today is Exchange Hosted Filtering, which provides spam filtering for on-premises Exchange mailboxes.

Jeff James: Some of our readers say that using AD is like going to the dentist—you know it's good for you and you know you need to do it, but it can be painful, from an ease-of-use perspective. How do your new products address those concerns, and how will they work with the new AD features in Windows Server 2008 R2?

Doug Leland: As you mentioned, Active Directory is the core, the heart and soul of any good identity infrastructure. Management of that system is key.

It's also consistent with what we're hearing from a customer-needs perspective of helping reduce the cost of these systems.

So that is an area we focused on for our 2008 release and are continuing to focus on for our upcoming release of Windows Server 2008 R2.

In terms of overall manageability, there are a number of significant advancements

that have taken place, and one of them is the adoption of PowerShell.

We are using PowerShell for all of our management interfaces, and that has dramatically increased the usability from an IT pro or administrative perspective.

We've also moved to a task-based paradigm. And within that paradigm, we can more easily identify and walk an admin through a particular task or a set of tasks if that's the way the interface is built up.

So, I think customers and administrators will see a huge benefit in terms of the overall manageability of the system.

In addition, we offer other products for managing identities and managing the life cycles of those identities and those resources in the organization.


One of those is Identity Lifecycle Manager, a tool that is designed to help organizations manage identities (users), manage groups, manage policies associated with those groups, and ultimately help them report on that and meet their compliance needs.

ILM 2007 is available for purchase today, and the next major release of that product, Identity Lifecycle Manager version 2, is currently in the release candidate (RC) phase.

Jeff James: Any estimate on when the final release of that might be?

Doug Leland: Well, the testing is going well—we released that RC back in November—and we're getting a lot of great feedback from customers.

We have a policy that you're probably familiar with, which is called dogfooding, and that is we won't release our enterprise products until we are running them in our own production environments.

We're working closely with MS IT in deploying that out, scaling that out right now actually, and we're moving towards the final release in a couple of months. 

InstantDoc ID 101758



Jeff James

(jjames@windowsitpro.com) is Editor-in-Chief, Web Content Strategist for Penton Media's IT Publishing Group. He specializes in server OSs, systems management, and server virtualization.

Check out the
best release of
the Windows
Server OS yet

by Michael Otey

Inside WINDOWS SERVER 2008 R2

Although it seems hard to believe, the next release of Microsoft's server OS, Windows Server 2008 R2, is right around the corner. At press time, Server 2008 R2 was in beta and scheduled for release late this year. Let's dive in and take a look at some of the most important new features in Server 2008 R2.

The Exclusive 64-Bit Club

Server 2008 R2 is the first Microsoft Windows Server OS to take the 64-bit-only road. This enhancement shouldn't be a problem for new installations, because most of today's servers are x64 compatible. However, Server 2008 R2 won't run on older 32-bit servers. Existing 32-bit applications can run on Server 2008 R2 by using the OS's 32-bit compatible Windows-On-Windows (WoW) subsystem.

In addition to going 64-bit-only, Server 2008 R2 also benefits from scalability enhancements. Server 2008 R2 can address as many as 256 logical processors on one server—up from a maximum of 64 in the original Windows Server 2008 release.

New Hyper-V Release

Another important enhancement in Server 2008 R2 is the inclusion of a new release of Hyper-V. A prerelease version of Hyper-V was shipped with the original Server 2008, then the final release was added as an update. The Server 2008 R2 version of Hyper-V can use more than 32 logical processors on the host virtual machine (VM). This new Hyper-V release can take advantage of the latest Intel and AMD Second Level Address Translation (SLAT) hardware virtualization support. SLAT lets the hypervisor dispense with shadow page tables and handles the translation of VM memory to physical memory, resulting in improved VM performance. Hyper-V in Server 2008 R2 also increases the memory support for VMs to 64GB. TCP offload and jumbo frames provide improved networking performance.

Another enhancement to Hyper-V in Server 2008 R2 is enhanced support for PowerShell management via a set of dedicated cmdlets. However, the single most important feature in Server 2008 R2 related to virtualization is support for Live Migration.

Live Migration

Live Migration is Microsoft's answer to VMware's VMotion. Live Migration lets you move Hyper-V VMs between Server 2008 R2 hosts with no downtime. Like VMotion, Live Migration lets the administrator handle planned downtime scenarios with no loss of VM availability. Live Migration requires Windows Failover Clustering and leverages Windows Clustering Services and the new Cluster Shared Volume (CSV) technology to move VMs between hosts in milliseconds. Server 2008 R2's new CSV technology lets multiple cluster nodes concurrently access the same LUN, which in turn lets them access the same Virtual Hard Disks (VHDs). Thus, the VHDs don't need to be physically moved to perform a Live Migration. Figure 1, page 24, shows an overview of how Live Migration works.

To perform a Live Migration, the administrator initiates the migration of a VM from the source node to a target cluster node. Live Migration creates a container VM on the target node. You don't need to move the VHD, because CSVs gives the target node full access to the VHD file stored on the SAN. Next, the source VM's current memory is copied to the target node. Clients connected to the source VM continue to run, and all the changed memory pages in the source VM are mirrored. The mirrored pages are then copied to the target VM until the delta is zero or until a finite number of iterations are reached. At that point Live Migration pauses the VM on the source, copies any remaining dirty pages, copies the partition state, starts the VM on the target node, and redirects all of the client connections from the source VM to the target VM. The migration is then complete and the source VM is deleted.




Figure 1: Live Migration overview

Active Directory Enhancements

From a Windows administrator's perspective, the biggest change in Server 2008 R2 is undoubtedly the new Active Directory Administrative Center. The ADAC provides a brand-new task-driven interface for managing Active Directory (AD). The older Users and Computers, Site and Services, and Domains and Trusts options still exist; however, the new ADAC's task-driven interface provides a better way to handle day-to-day tasks such as working with users, computers, groups, and organizational units (OUs). The ADAC is installed when you run Dcpromo to make a Server 2008 R2 system a domain controller (DC). Figure 2 shows the new ADAC.

As you can see in Figure 2, the ADAC provides breadcrumb-style navigation much like Windows Explorer. By default, the ADAC navigation pane on the left side of the screen uses either a treeview or a simple list view. However, you can also customize the view by adding commonly used containers to the navigation pane. The new ADAC can open AD using a different set of credentials than your logon credentials. It can also manage AD objects across multiple domains. The current version of the ADAC runs only on Server 2008 R2. Future versions of Windows 7 will also be able to run the ADAC. For more information about the ADAC, see the Microsoft TechNet article "What's New in AD DS: Active Directory Administrative Center," at technet.microsoft.com/en-us/

library/dd378856.aspx.

A closely related AD enhancement in Server 2008 R2 is the addition of 75 new AD cmdlets. The new ADAC is built using these cmdlets. When you use ADAC to perform actions, ADAC is actually building and executing PowerShell scripts in the background.

The new Server 2008 R2 Active Directory Domain Services also has several significant improvements. A new offline domain join feature lets a computer join a domain without being connected to the domain. This feature can help automate client deployment. Another useful feature in the new Active Directory Domain Services is the new AD Recycle Bin, which lets you recover deleted AD objects without performing an authoritative restore.

Remote System Management with Server Manager

One feature that administrators love in Server 2008 is Server Manager. Server Manager provides a centralized management console that is actually useful. You can use Server Manager to manage roles and features, as well as check status and drill into event logs. However, in the original Server 2008 release Server Manager is limited to working with the local system. The Server 2008 R2 release enables Server Manager to manage both local and remote Server 2008 systems. In

addition, Server Manager can be installed on Windows Vista or Windows 7 network clients, letting you perform network management tasks from client workstations.

Terminal Services Is Out; Remote Desktop Services Is In

Another change in Server 2008 R2 is the rebranding of Terminal Services to Remote Desktop Services. Web Table 1 (www.windowsitpro.com, InstantDoc ID 101706) lists the former names of various Terminal Services components and their corresponding Remote Desktop Services names.

Server 2008 R2's Remote Desktop Services changes aren't just in name alone. The new RemoteApp & Desktop Connection (RAD) includes support for the Aero Glass interface, true multi-monitor support, multimedia redirection, audio recording, and support for DirectX 9, 10, and 11 redirection.

Enhanced Scripting Functionality with PowerShell 2.0

Server 2008 R2 includes the new PowerShell 2.0 release. PowerShell 2.0 is compatible with PowerShell 1.0; has improved Windows Management Instrumentation (WMI) cmdlets; and supports running scripts on remote systems, creating Script Cmdlets, and running background jobs. More than 240 new cmdlets ship with Server 2008 R2 out of the box.

Even better, Server 2008 R2 provides a new graphical PowerShell UI for developing

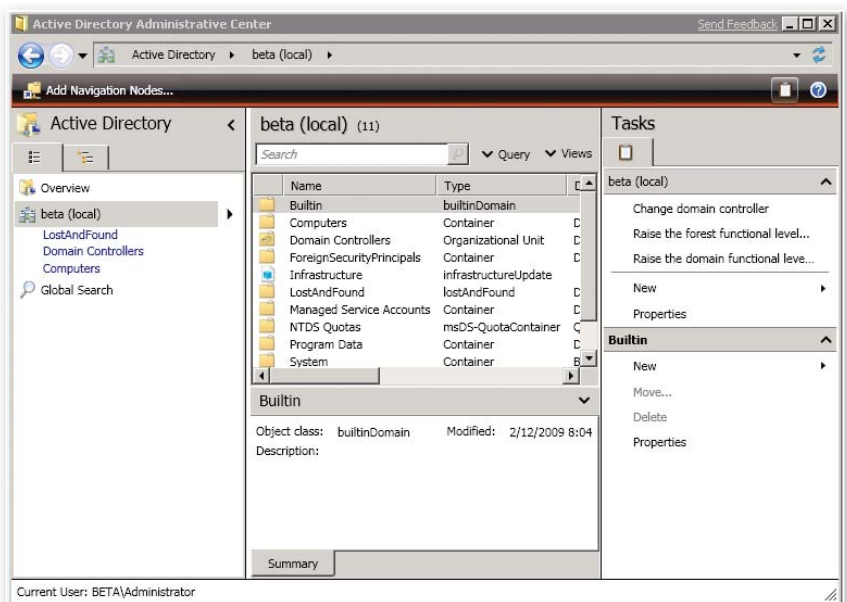


Figure 2: Active Directory Administrative Center

and debugging PowerShell scripts. The new PowerShell Integrated Scripting Environment (ISE) is a multi-tabbed graphical PowerShell development platform. The PowerShell ISE features color-coded syntax, as well as debugging capabilities that let you set breakpoints and single-step through your PowerShell scripts. Figure 3 shows the new PowerShell ISE.

The PowerShell ISE development environment consists of three panes: the script pane, the output pane, and the command pane. You can see the script pane in the upper third of Figure 3; this pane is for editing and debugging your PowerShell scripts. The output pane (in the middle portion of Figure 3) displays the results of any scripts that you can run in the ISE. The command pane, which you can see in the bottom part of Figure 3, is for executing your scripts and other PowerShell commands.

.NET Framework Support in Server Core

One of the biggest disappointments in the original Server 2008 release was the lack of support for the .NET Framework in Server Core. Several technologies that seemed perfect for Server Core, such as PowerShell and ASP.NET applications, couldn't run on Server Core. Server 2008 R2 adds support for the .NET Framework versions

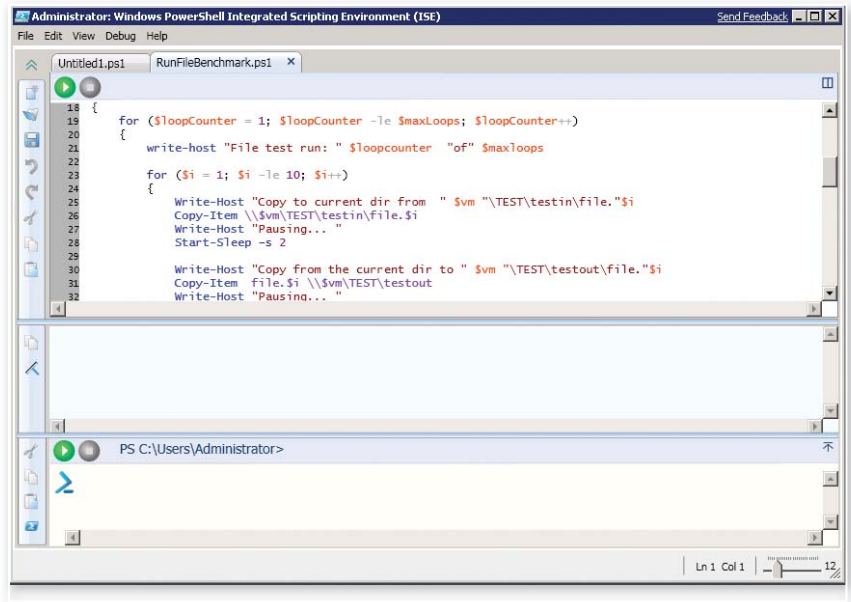


Figure 3: PowerShell Integrated Scripting Environment

2.0, 3.0, 3.5, and 4.0. Support for the .NET Framework allows Server Core to run both ASP.NET applications and PowerShell scripts. However, Server Core 2008 R2 still doesn't support SQL Server or Exchange.

IIS 7.5

Another new feature in Server 2008 R2 is the inclusion of IIS 7.5. The main enhancements in IIS 7.5 are improved management and deployment of web applications. IIS

7.5 has a new PowerShell Provider for IIS, along with several new IIS task-oriented PowerShell management cmdlets. The new cmdlets provide the ability to add and change configuration properties of websites, web-based applications, virtual directories, and application pools.

IIS Manager is also enhanced with a new Configuration Editor. This feature lets you access all of the IIS 7.5 configuration settings, including settings such as FastCGI that

BID2LEARN.COM™

Bid. Save. Learn.™

Never pay retail again for training Classes!

Log On and book that class!

Our Place Your Bid™ Auction Service is a revolutionary, service, providing the consumer an opportunity to reduce the cost of a class in a "hassle free" negotiated environment; much the way the travel and lodging arena have.

- Classes
- Courses
- Licensing
- Bootcamps
- Certifications

Same National Vendors & Local Trainers

you already use!

IT Art Business Colleges Culinary Healthcare Vocational Recreational

www.Bid2Learn.com™

were hidden in the previous version of IIS.

For added web application security, a new Request Filtering module provides HTTP blocking capabilities that were formerly found in the separate URLScan product. IIS 7.5 also provides improved application security by running every application pool with a unique low-privilege identity. Also included are a new Best Practices Analyzer (BPA) and updated versions of Secure FTP and WebDAV.

Going Green with Core Parking and P-states

Server 2008 R2's Core Parking feature lets the OS dynamically control the number of cores used in a multi-core server. Core Parking continually monitors the CPU utilization of multi-core server systems. When ever processor cores are underutilized, Core Parking can put those cores into sleep mode to reduce the power required to run the system. When the workload on the remaining cores increases, the suspended cores are reactivated and full processing power returns. For example, Core Parking

could enable a server with 64 logical cores to drop back to just a 2-core machine during low-utilization times, then restore the server to a full 64-core system when the workload rises. Notably, with Core Parking one core must always be active in order to control the state of the other cores.

Another power-management feature built into Server 2008 R2 is the ability to adjust processors' Advanced Configuration and Power Interface (ACPI) P-states. This feature essentially allows very granular control over a system's power consumption. Altering the P-state of the processor governs the frequency of the CPU. Running the processor cores at lower frequencies is another way to reduce power consumption. Both Core Parking and ACPI P-state status can be controlled through new Group Policy settings.

Best Windows Server OS Yet

Server 2008 R2 adds a lot of value to the Server 2008 OS, with features such as Live Migration, the new ADAC, and the PowerShell ISE. Other important features include

enhanced DNS and DHCP security, read-only DFS Replication (DFSR), and the ability to boot from VHD. Connecting Windows 7 to Server 2008 R2 provides even more benefits; for more information, see the web sidebar "Windows Server 2008 R2 and Windows 7: Better Together," InstantDoc ID 101707. In addition to the big-ticket items, Server 2008 R2 provides numerous smaller changes. For an overview of these changes, see the web sidebar "The Little Things About Windows Server 2008 R2," InstantDoc ID 101708.

I used the beta version to evaluate Server 2008 R2; some features will likely change before the final release. However, Server 2008 R2 is clearly the best release of the Windows Server OS yet.

InstantDoc ID 101706



Michael Otey

(motey@windowsitpro.com) is technical director for *Windows IT Pro* and *SQL Server Magazine* and author of *Microsoft SQL Server 2008 New Features* (Osborne/McGraw-Hill).

Next-Generation Enterprise Backups

R1Soft Continuous Data Protection®, Tools for Windows & Linux

****Maximum Protection at a Lower Cost Per Server****

Key Features	R1Soft CDP Server	Acronis® True Image	IBM Tivoli®
Daily Backups			
Hourly Backups		Not Supported	Desktops Only
Bare-Metal Restore			Limited
Continuous Data Protection® CDP		Not Supported	Desktops Only
CDP for MySQL		Not Supported	Not Supported
Easy-to-use Web Interface		Not Supported	Not Supported
MS VSS Integration			
Support Xen & VMWare			
	\$180/server	\$699 /server	You Can't Afford It

CDP for MySQL Add-on: True-Granular Restore™, Bare-Metal Restore for MySQL databases, Store more than 100 recovery points per day, Table Level Restore for MySQL, Zero Downtime during Backups and more...

For more information visit: www.r1soft.com



R1Soft

Continuous Data Protection®
Tools for Windows and Linux

R1Soft, a division of BBS Technologies, Inc. 2009 © Copyright BBS Technologies, Inc.

Organize Your Active Directory Objects

Do you have a good handle on which objects in your directory are no longer used? Do you know exactly who you need to contact when making changes to the content or structure of your forest? As a consultant specializing in Active Directory (AD), I come across many

AD implementations that have grown organically over time. Typically, these implementations contain a large number of unused objects, as well as objects that are obviously in use, but who or what is using them isn't clear. It's costly having objects in this state: Periodic cleanups of AD become labor intensive and expensive, AD restructures or migrations become more complex, and even simple change management is more difficult.

To gain control over your AD environment, you need to deal with three key elements of object lifecycle management. The first is determining the appropriate way to provision, re-provision, and de-provision objects. The second is setting up controls so that all new objects conform to the provisioning methodology. The third is the sometimes arduous and time-consuming work of cleaning up existing objects so they either conform to the methodology or can be deleted from AD. In this article I provide advice and tips that will assist you with the first two aspects by introducing the concept of "guardianship" of AD objects. By associating real people ("guardians") with AD objects, you can gain greater control over your AD environment. I also offer some examples to help with the third aspect, clean-up.

Clarifying Terminology

The guardian for an AD object is the human being directly responsible for, or most closely associated with, that object. A better term might be "contact," but I'll avoid that because it's already used to represent a specific type of object in AD. Another term might be "owner," but this too has meaning in AD security in the context of the creator/owner of an object.

It would be handy if there were an AD attribute named "guardian" that we could use for setting guardianship of different types of AD objects. Unfortunately, there isn't so we must either create a new attribute (which involves extending the schema), or use an existing attribute from the default AD schema. For simplicity and because most organizations

"Guardian"
concept helps
to control object
lifecycle

by Tony Murray

■ ORGANIZE AD OBJECTS

have a healthy aversion to extending the schema, I use existing attributes as you'll see in the following sections.

Benefits of Guardianship

Identifying and removing unused objects in AD can be a thankless and time-consuming task. You can find tools to assist you with locating unused objects (the Windows command-line tool `dsquery` is one; `AdFind` and `Old-Cmp` from www.joeware.net are others), but because object deletion is potentially damaging to systems and applications that leverage AD you need to be 100 percent sure that you're dealing with an unused object before you delete it. Typically you would check with the person currently responsible for that object. But in many cases this person isn't easily identifiable from the object's attributes. You might have only the object name to work with (e.g., a group named "OKP100 Staff"). This is fine if OKP100 means something to you, but otherwise it's no help at all. The object's description might contain some information (e.g., "See JP Carter before making changes"), but what if JP Carter no longer works for the organization? As you can see, no magic built-in feature automatically links a human owner to an AD object: It's something that you have to implement for yourself. This is where the guardianship concept can help you.

The guardianship concept can also assist you when working with active objects. For example, when processing a request to add a user to a group, your operational staff can refer to the guardian to approve or decline the request. The suggestions I make for setting guardianship of objects assume that you will use AD as the repository for guardianship information. The same concepts (but clearly different methodology) apply if you already have a tool in place for provisioning AD objects and that tool is capable of storing the required guardianship information.

Setting Guardianship for User Objects

Organizations create user objects for a range of different purposes. Aside from standard

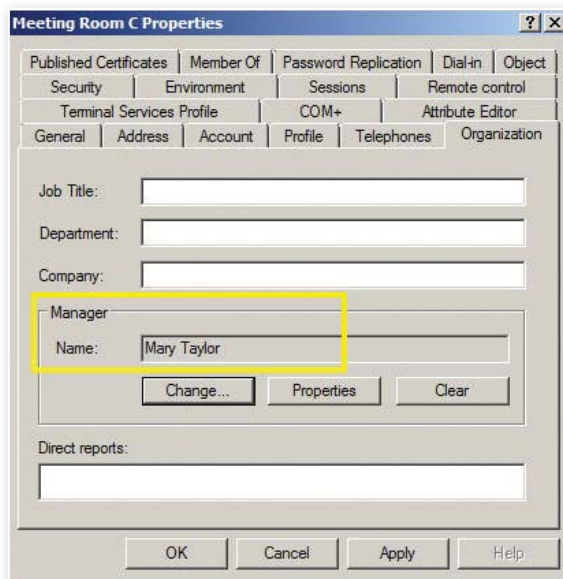


Figure 1: Setting up a guardian for a resource account

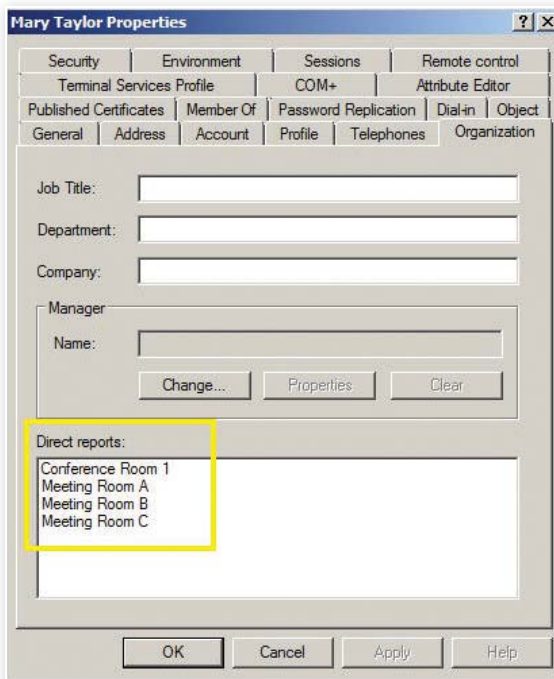


Figure 2: Using a linked attribute to set up a resource account as a direct report

user accounts directly associated with a warm body, user objects can be created for shared accounts, resource accounts (for mailboxes such as meeting rooms), service accounts, and secondary accounts for administrative purposes. For all types of user objects, I recommend associating a guardian by setting the value of the manager attribute. Let's look at an example in which we have a resource account for a meeting-room mailbox named Meeting Room C. We want to set the guardian to be Mary Taylor.

From within the Active Directory Users and Computers MMC snap-in, find Meeting Room C, open up the properties and select the Organization tab. From here, click Change within the Manager section and use the object picker to find and add Mary Taylor's user account, as Figure 1 shows.

The manager attribute is a linked attribute. (For more information, see "Linked Attributes," at [msdn.microsoft.com/en-us/library/ms677270\(VS.85\).aspx](http://msdn.microsoft.com/en-us/library/ms677270(VS.85).aspx).) The manager attribute is the forward link, while `directReports` is the corresponding back-link attribute. Because the attributes are linked, when I set Mary Taylor as Meeting Room C's manager, Mary Taylor's user object shows Meeting Room C as a direct report, which Figure 2 shows. The main advantage of using a linked attribute is that the link object can be renamed or moved within AD, and the link remains intact. The link can only be broken if either the forward or back-link object is deleted. Another advantage of the linked attribute is that it lets you search AD for the relationship using either the guardian or the object(s) for which the guardian is responsible.

Following are examples of such searches using the `AdFind` tool from www.joeware.net. The first example shows a search for all user accounts for which Mary Taylor is the guardian:

```
C:\>adfind -list -b "CN=Mary Taylor,OU=Standard User Accounts,DC=ad,DC=fisheagle,DC=net" directReports
```

Figure 3 shows the results of that search.

The second example shows a search for a meeting room's guardian:

```
C:\>adfind -list -b "CN=Meeting Room C,OU=Resource Accounts,DC=ad,DC=fisheagle,DC=net" manager
```


The results of that second search show Mary Taylor as the guardian: CN=Mary Taylor,OU=Standard UserAccounts,DC=ad,DC=fisheagle,DC=net.

Setting Guardianship for Group Objects

The manager and directReports linked-attribute pair isn't available for use with groups. Instead, I recommend using a similar pair of linked attributes named managedBy and managedObjects.

Let's look at an example in which we have a group named Consulting Team. We want to set the guardian to be Mary Taylor. To do this, locate the group within Active Directory Users and Computers, open the properties, and select the Managed By tab. From here, click Change within the Name section and use the object picker to find and add Mary Taylor's user account, which you can see in Figure 4.

When you make the change in Active Directory Users and Computers, AD sets the value of managedBy on the group object as the distinguished name (DN) of Mary Taylor's account. Note that when setting the managedBy value, you have the option to select *Manager can update membership list*, as Figure 4 shows. If you need only to assign guardianship for informational purposes, then you probably don't want to select the option, but it otherwise provides a shortcut method of assigning delegated management of the group membership to the guardian.

You can then query AD by using the object for which the guardian is responsible. The AdFind example below shows the managedBy value in use:

```
C:\>adfind -list -b "CN=Consulting Team, OU=Groups,DC=ad,DC=fisheagle,DC=net" managedBy
```

Here are the results of that search: CN=Mary Taylor,OU=Standard User Accounts,DC=ad,DC=fisheagle,DC=net.

You can also search for the relationship between the group and its guardian by querying AD using the guardian:

```
C:\>adfind -list -b "CN=Mary Taylor,OU=Standard User Accounts,DC=ad,DC=fisheagle,DC=net" managedObjects
```

```
CN=Meeting Room C,OU=Resource Accounts,DC=ad,DC=fisheagle,DC=net
CN=Meeting Room B,OU=Resource Accounts,DC=ad,DC=fisheagle,DC=net
CN=Meeting Room A,OU=Resource Accounts,DC=ad,DC=fisheagle,DC=net
CN=Conference Room 1,OU=Resource Accounts,DC=ad,DC=fisheagle,DC=net
```

Figure 3: Search results showing all of a guardian's user accounts

again using AdFind for the example. Figure 5 shows the results.

Be aware that the managedObjects backlink isn't visible in Active Directory Users and Computers in the way that directReports is. To view the back link, you need to use LDAP queries or tools such as ADSIEdit or the new Attribute Editor in Windows Server 2008.

Setting Guardianship for Other Object Types

You can extend the concept of ownership to include any other type of AD object. For example, computer and organizational unit (OU) objects spring to mind as likely candidates. Both of these support the use of the managedBy and managedObjects linked attributes, so you can use the same method as for groups to define the guardian relationship.

Guardianship and Object Lifecycle Management

The guardianship concept works best if it's incorporated into your organization's provisioning and de-provisioning procedures. It's important to set the guardian whenever you provision an AD object that you want to keep track of. Similarly, when you de-provision a standard user account, you should ensure that any guardianship relationships associated with that user are either removed or transferred to another user account. For example, if Mary Taylor from our scenario leaves the company, I would need to consider what to do with the objects

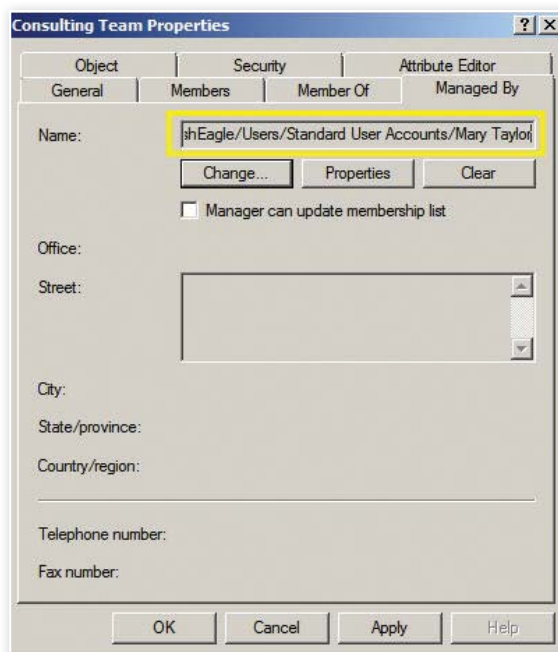


Figure 4: Setting up a guardian for a group

```
CN=Management Team,OU=Groups,DC=ad,DC=fisheagle,DC=net
CN=Project Management Team,OU=Groups,DC=ad,DC=fisheagle,DC=net
CN=Consulting Team,OU=Groups,DC=ad,DC=fisheagle,DC=net
```

Figure 5: Results from querying AD using the guardian

for which she is guardian. In the case of groups and meeting rooms, I would probably transfer the guardianship to another account. If Mary has a secondary account for administrative purposes for which she is guardian, I would probably de-provision that at the same time.

Also bear in mind that people often change roles within an organization. When this happens, your re-provisioning procedures should reflect the fact that someone might no longer be the appropriate guardian for AD objects.

Cleaning Up Your Existing AD Infrastructure

With a working guardianship in place for newly provisioned objects, you need to address the task of identifying existing objects that have no guardian, then

■ ORGANIZE AD OBJECTS

configure the appropriate guardian relationships. Before you do that, however, you should remove any objects from AD that are no longer required.

Within AD, objects often are created for a specific reason and simply forgotten. For example, accounts and groups might be created to support a new document-management system. If the organization decommissions the document-management system a few years down the line, the associated accounts and groups might not

This attribute also lets you detect stale objects by querying a single domain controller (DC) rather than attempting to consolidate the lastLogon results from all DCs in the domain. The lastLogonTimestamp attribute is available with Windows Server 2003 and later.

Finding Inactive Computer Objects

Similar to the search for inactive user objects, the following example, written as one line, uses AdFind to search for computer objects

confirmation isn't received within XX days, you could then initiate the de-provisioning process for that group.

In the following example, I use AdFind to search for groups with no members, which is one indicator that the group might not be in use. Note that I exclude critical system objects (e.g., Enterprise Admins, built-in groups) as these can be legitimately empty and should never be removed:

```
adfind -csv -default -f "(&(objectcategory=group)(!member=*)(!isCriticalSystemObject=TRUE))" samaccountname description managedby
```

A lack of members might indicate an unused group, but realistically the only reliable method to determine whether a group is still in use is to set up a guardian for it.

be removed in parallel. These stale objects could linger in AD indefinitely or until someone questions their existence.

In the next two sections, I use the AdFind command-line tool to find unused objects. Note that I could instead have used OldCMP or the built-in dsquery tool for the inactive user and computer object searches.

Finding Inactive User Objects

The following example, written all on one line, uses AdFind to search for user accounts that have either never logged on to the domain or haven't logged on to the domain since the beginning of 2008. The output is in CSV format.

```
adfind -csv -default -tdca -utc -binenc -bit -f "(&(samaccounttype=805306368)(!(lastLogonTimestamp<={utc:2008/01/01}))(!(lastLogonTimestamp=*)))(!(userAccountControl:AND:=2)))" lastlogontimestamp pwdlastset accountexpires whencreated
```

The search excludes disabled users because most organizations tend to leave de-provisioned user objects in a disabled state for a period of time prior to deleting them. Additionally, the search uses the lastLogonTimestamp attribute, a replicated attribute that gets updated periodically (and which is consequently not as accurate as the non-replicated lastLogon attribute).

that have either never logged on to the domain or have not logged on to the domain since the beginning of 2008:

```
adfind -csv -default -tdca -utc -binenc -f "(&(objectcategory=computer)(!(lastLogonTimestamp<={utc:2008/01/01}))(!(lastLogonTimestamp=*)))))" name operatingSystem operatingSystemServicePack lastlogontimestamp pwdlastset whencreated
```

Finding Unused Groups

It's very difficult to determine whether a group is still in use within AD. At least with user and computer objects you can query for when the user or computer last set the password (using the pwdLastSet attribute) or when the user or computer account last logged on (using the lastLogonTimestamp attribute).


But groups don't have passwords and don't log on to AD, so there are no useful attributes to help you determine whether a group is still in use. A lack of members might indicate an unused group, but realistically the only reliable method to determine whether a group is still in use is to set up a guardian for it.

You could then set up a process periodically requesting confirmation of the guardians that a group is still in use. If a

It's important to test the validity of your search results. An LDAP search against AD for the information is just one aspect of the overall task. You should qualitatively assess each result. For example, it might be entirely valid for a user object corresponding to a resource mailbox (e.g., a meeting room) not to have logged on to the domain for 12 months or more. Another example is a group that has no members but is required to be present for a specific application to function.

Minimal Effort, Maximum Reward

Whatever terminology you use—manager, owner, contact, or guardian—the concept of linking AD objects to real people isn't new. In fact, Microsoft makes some provision in AD for defining the relationship through the managedBy and managedObjects linked-attribute pair for use with certain object types.

I strongly recommend that you consider implementing the concept of guardianship in your environment. The effort involved in setting up the required procedures is low and far outweighs the cost of dealing with an uncontrolled environment. The sooner you do this, the less effort you will need to spend on clean-up tasks at a later date. 

InstantDoc ID 101655



Tony Murray

(tony@fish-eagle.net) works for Fish Eagle Limited, a messaging and directory service consultancy in New Zealand, and runs the ActiveDir.org website. He is an MVP (Directory Services) and an MCSE.

A First Look at Exchange 14

2010

Anyone dedicated to trivia will note that the code name for Microsoft Exchange Server 2007 was Exchange 12, but the next major release of its mail server has been code named Exchange 14. Microsoft skipped 13 for the same reason that many hotels don't have a thirteenth floor—superstition!

Exchange 14 is expected to ship in late 2009 and have a final name of Microsoft Exchange Server 2010. Exchange 2010 follows up the architectural changes made in Exchange 2007 with some big updates of its own to give the product better performance and make it more resilient and easier to manage. The most important changes fall broadly into the categories of an Information Store refresh, a new approach to high availability, management and administration updates, and messaging compliance improvements.

Enhancements to the Store

Exchange has always been a challenging application for storage because the I/O profile of a busy mailbox server consists of many random small I/O operations rather than the predictable I/O patterns you see in other database-centric applications. This situation can be explained by the huge variety of messages that an Exchange server handles—from the simple, one-line message sent to a single recipient to the multimegabyte message (including attachments) sent to nested distribution lists. Obviously these transactions create radically different I/O demands.

Microsoft greatly reduced disk I/O with Exchange 2007, largely by trading the extra memory made available by using the 64-bit platform to cache as much Store data as possible. This process resulted in a significant I/O reduction per active mailbox—except in the case of large mailboxes. The problem with large mailboxes is that users tend to keep thousands of items scattered around hundreds of folders. The more items and folders in a mailbox, the more work the Store has to do to organize and maintain the indexes that underpin the mailbox. Windows Desktop Search with its Microsoft Office Outlook integration lets users become even less organized: If they forget where something is in their large folder structure, it's easy to perform a search to find the desired item.

So, although Exchange 2007 made real improvements by optimizing Store caching, human behavior meant that further work was necessary for Exchange to effectively support very large mailboxes. As it happens, Microsoft had previously assessed whether they could

Microsoft's latest mail server introduces improved high availability and easier management

by Tony Redmond

move the underlying Store database engine from Extensible Storage Engine (ESE) to Microsoft SQL Server. The engineering investment to make this change proved too great, which is why Exchange still uses ESE. However, the investigation reviewed some fundamental aspects of the Store database, including its schema and tables. As a result, some changes to aid performance were included in Exchange 2007, notably the increase in page size from 4KB to 8KB and smoother I/O transactions. Further performance improvements in Exchange 2010 include:

- Increased page size from 8KB to 32KB—With this change, more data can be stored in a single page, avoiding the need to scatter across the database the pages required for a single item, including any attachments.
- Header data for all mailbox items is stored in a single database table—This change makes the database more efficient because it can process a single table for a mailbox during a client session instead of accessing different tables for different mailbox folders. A side effect of this schema change is that Exchange no longer uses Single Instance Storage (SIS) to keep just one copy of message content per database. Most servers support multiple databases, so the efficiency gained from SIS is less and less as time goes on.
- The Store compresses attachments—Microsoft calculates that the CPU time spent compressing and decompressing attachments is less than the work required to manage very large uncompressed data within the database. This change also reduces the overall size of Exchange databases, which speeds up operations such as backups.
- The Store updates views (indexes) only when they're accessed—An Outlook client can create many different views for a folder on the fly (e.g., items ordered by subject), and the Store maintains these views within the database. The Store ages out unused views after 40 days, but it needs to maintain views until then. Updating views only when needed eliminates a lot of background processing.

Microsoft's initial performance results indicate that the new Store generates

substantially fewer I/O operations than its Exchange 2007 equivalent. Reducing I/O lets servers support more mailboxes as well as allowing additional flexibility in storage options. Traditionally, large mailbox servers have used high-end storage configurations such as SANs to deliver excellent I/O performance with maximum reliability. If Exchange 2010 delivers a smaller I/O footprint and better resilience, system designers might be tempted to use lower-cost Serial ATA (SATA) and Just a Bunch of Disks (JBOD) storage. Changes will still occur in the code before Exchange 2010 ships, so we'll have to wait a bit to know how to optimize storage for production environments.

High Availability at the Core

Exchange 2007 introduced log shipping to let administrators replicate data to local disks (local continuous replication—LCR), to another node in a cluster (cluster continuous replication—CCR), and to a server in another data center (standby continuous replication—SCR). Microsoft builds off this log shipping technology to make high availability a core characteristic of Exchange 2010. Microsoft is shaking up Exchange's high availability feature set through four key steps:

- The concept of storage groups is eliminated, so the database becomes the management unit to plan high availability around—this is a sensible step given that log replication works only for a storage group containing a single database.
- Single copy clusters are eliminated and not supported in Exchange 2010. Microsoft is moving toward the idea that maintaining multiple copies of data on multiple servers delivers better high availability than attempting to update a single copy of data. Microsoft has also removed LCR from Exchange 2010 because log replication on the same server delivers limited value.
- Exchange 2010 introduces Database Availability Groups (DAGs), which are groupings of up to 16 servers in which some or all of the databases are marked for replication to one or more other servers. Microsoft uses some components of Windows clustering (e.g., heartbeats, the file share witness) to connect servers within the DAG, which can span physical locations. The big feature is that you

can replicate databases to multiple servers within the DAG through log shipping, so locations within a DAG must share sufficient network resources to be able to copy logs quickly enough so that queues of unplayed logs don't build up; think of this requirement as being similar to that of SCR today. Replication targets are chosen at the database level rather than the server level, so you can replicate different databases from a server to different servers within the DAG. For example, a server in New York that has two databases could replicate one database to a server in Los Angeles and the other to a server in Seattle. The live database is referred to as the master; if a problem occurs with the master database, a component called the Active Manager switches to one of its replicas and makes it the live master. Microsoft includes management for DAGs in Exchange 2010's version of Exchange Management Console (EMC) and adds Exchange Management Shell (EMS) commands, so you can control DAGs through the GUI or the command line.

- A new component in Exchange 2010 called the RPC Client Access Layer upgrades the Client Access server role so that all client connections flow through a predictable point in the network. With the potential for live copies of databases switching between servers, clients can become confused when they attempt to connect to a mailbox. Exchange 2007 introduced the Client Access role, which manages connections from all clients except MAPI (i.e., Outlook). In Exchange 2010, the Client Access role determines which server currently hosts the live copy of a mailbox by reference to the DAG information, which is held in Active Directory (AD), and is therefore able to redirect clients when a database has been switched.

There are challenges with any high availability solution. Some obvious problems that deserve consideration are how third-party backup software will deal with DAGs and what role offline backups play after you deploy Exchange 2010. The introduction of DAGs indicates that Microsoft is heading toward multiple database replicas as the primary solution for data availability. Because

May 2009

The Essential Guide to **Exchange Server Consolidation**

By Paul Robichaux

UNISYS

SPECIAL ADVERTISING SUPPLEMENT TO *WINDOWS IT PRO*

Microsoft has steadily improved the scalability, availability and performance of Exchange Server since its initial release. In many environments, Microsoft's standard recommendations for sizing and configuring servers work perfectly well. However, there are several additional ways in which you can optimize Exchange architecture designs to reduce cost and increase service levels. One method is to deploy multiple Exchange Server roles onto a smaller number of servers, sometimes using more powerful physical hardware, with the aid of hardware virtualization solutions. Another is to put more mailboxes on each mailbox server instance, although this requires exceeding the (admittedly conservative) recommendations and best practices now used for Exchange 2007 deployments.

You might wonder how a consolidated Exchange deployment is different from an ordinary deployment. Unisys has developed a packaged consulting offering called Consolidated Exchange Solution (CES) that codifies the differences in deployment and scalability. Even if you're not using CES, though, you may still benefit from applying the CES principles to your own deployment.

Benefits of Consolidation

Modern IT has seen a long-running battle between consolidation and decentralization. Starting in the era of the mainframe, there's been a tension between the business benefits of maintaining centralized IT resources managed by a staff of experts who allocate resources and the benefits of letting individual organizational units, teams, or people have their own resources to use as they see fit. Periodically, technological changes shift the balance one way or another. Right now, the pendulum is swinging strongly towards consolidation because the potential benefits are too good to ignore:

- Improved cost control: consolidating gives you more predictability in costs through standardization and elimination of redundant hardware and software licenses.
- Reduced capital expenditure for new asset requirements: there's always tension between providing adequate room for growth (which requires you to spend money now to buy capacity you won't immediately need) and having to add capacity on an ad hoc basis (which is operationally disruptive and hard to plan).

- Better operating efficiency: a properly implemented consolidation will help you get the most possible use from your existing hardware. Consider the number of servers in your environment that loaf along using 25 percent or less of their CPU—they're probably great candidates for consolidation.
- Power and cooling savings: the fewer servers you have, the less money you'll have to spend to power them, and the less cooling your data center will require. Data center power consumption often represents a hidden expense because IT operations typically don't pay for their own power—some other part of the parent company does. However, lowering power and cooling bills is a great way to attract positive attention for your company; witness the efforts by IBM, Google, Sun, and others to showcase their "green data center" solutions.
- Better business continuance support: for most highly decentralized environments, it's not feasible to deploy business continuance or HA solutions like stretched CCR clusters. The individual deployments aren't that expensive, but when you multiply those individual costs by the number of deployments required, the numbers frequently don't make sense. Centralized environments allow much more cost-effective HA and continuance deployments.
- Reduced infrastructure complexity: the fewer physical machines you have, the easier your environment will be to understand, maintain, and troubleshoot. This is particularly true if you can consolidate your DNS and Active Directory infrastructure as part of your overall Exchange consolidation plan.
- Reduced data center floor space: not only do servers turn power into heat, but they also take up room! Adding more power or cooling capacity to an existing space can be very expensive, but adding more space itself can be even more costly.

Combining Consolidation with Virtualization

Traditional Microsoft-style consolidation revolves around consolidating work. For example, the performance gains provided by Exchange 2007's use of 64-bit RAM address spaces to provide massive caches means that servers can handle many more concurrent mailboxes than they could with Exchange 2003; in the same vein, Exchange 2007's disk I/O improvements deliver the ability to host more and larger mailboxes on a given storage system than previous versions. Using larger servers to host more

Consider Virtualization as an Option when Consolidating Exchange Servers

Q How does a consolidated and virtualized environment differ from a more traditional Exchange deployment?

A The usual method of deploying Exchange is to size a server for a given number of mailboxes, then deploy as many “cookie cutter” servers as needed to handle the total load. This approach offers good standardization, but it can result in higher-than-necessary server counts. Combining infrastructure servers by virtualizing them and putting them on a larger physical server is one popular approach; another is to virtualize mailbox servers too, using SAN storage for the greatest possible efficiency.

Q Which is more useful: consolidation or virtualization?

A Both have their place! Consolidation refers to moving more work onto a smaller number of servers, usually by increasing the number of mailboxes on a given mailbox server. Virtualization means using a tool like Microsoft’s Hyper-V or VMware’s ESX to turn physical servers into virtual servers. You can combine these technologies to great effect with Exchange Server 2007. There are several ways in which you can optimize Exchange environments, resulting in reduced cost and increased service levels. One method is to deploy multiple Exchange Server roles onto fewer, and sometimes larger, physical servers with the aid of hardware virtualization platforms. Another is to deploy more mailboxes per mailbox server, reducing the number of servers needed to support a given workload.

Q Does Microsoft support virtualization of Exchange Server 2007?

A Yes. However, you need to be running the 64-bit production version of Exchange Server 2007 on Windows Server 2008, and you have to use a supported virtualization tool. At present, that means either Microsoft’s Hyper-V or VMware’s ESX 3.5, although other environments may be added in the future. Microsoft’s full support statement for virtualized Exchange Server 2007 can be found at <http://technet.microsoft.com/en-us/library/cc794548.aspx>.

Q What server roles can I effectively consolidate?

A Almost all of the Exchange Server 2007 server roles can be consolidated, either through physical consolidation or virtualization. The Unified Messaging role is not supported by Microsoft in a virtualized environment, but it can be consolidated, as can all the other Exchange roles.



Ted Buiting
Director of Communications and
Collaboration Solutions, Unisys

Ted Buiting is the Director of Communications and Collaboration Solutions for Unisys, for whom he has worked for over 25 years in a variety of finance and marketing roles. Ted has been a primary developer of solutions that enable organizations to make most efficient use of communications technology to advance the goals of their business. He has spearheaded efforts to demonstrate the significant operational and related ROI benefits, as well as the limitations, associated with the disciplined deployment of consolidation and virtualization.



users “per box” is an effective approach in many environments because it lowers the overall server count, which leads to the benefits described above.

Virtualization technology has become an important part of messaging architects’ toolkits. That’s because it provides an alternative to increasing the concentration of work assigned to a single physical server. The traditional deployment methodology is to size a server to handle a certain number of mailboxes, then deploy the correct number of servers to handle the expected number of mailboxes. Virtualization turns that model sideways with a focus on getting the best possible utilization from your servers by combining multiple virtual workloads onto each physical host.

What about combining virtualization and consolidation? It turns out these two deployment tools go together like rice and gravy. By virtualizing some or all of your servers, then running those virtualized servers on a small number of very powerful physical host machines, you can gain some stunning reductions in cost. Unisys claims server count reductions of up to 90 percent in some cases, with corresponding cost reductions in hardware maintenance (up to 80 percent), power and cooling (up to 70 percent), and connectivity and networking (up to 80 percent).

Microsoft and Virtualization

What does Microsoft think about virtualizing Exchange? There are two aspects to consider. First is what Microsoft actually supports, and the second is what Microsoft recommends.

The support policy is fully explained on Microsoft’s web site (see <http://technet.microsoft.com/en-us/library/cc794548.aspx>). To briefly summarize, Microsoft will support the 64-bit version of Exchange 2007 SP1 when it runs on Windows Server 2008 under Hyper-V or any virtualization solution that’s been certified under the Windows Server Virtualization Validation Program (SVVP). You have to dig around the SVVP site a bit before you find that VMware ESX 3.5 is certified. As with the clustering hardware compatibility list (HCL) of old, specific combinations of CPU types, cores, and RAM are certified. Microsoft doesn’t support using some virtualization features (including dynamic and differencing disks), and the unified messaging (UM) server role is not supported as a virtualization target. There are also some restrictions on what the physical host can do. For example, running applications like Exchange or SQL Server on the physical host is unsupported.

As for Microsoft’s recommendations (contained at the same URL), they’re too long to summarize

here, but the first two sentences of the recommendations make a great summary: “Running Exchange 2007 SP1 in a guest virtual machine does not change the Exchange Server design requirements from an application perspective. The Exchange Server guest virtual machine must still be sized appropriately to handle the workload.” Simply put, this means that your consolidated environment has to meet the same performance levels that a physical server would.

Taking Advantage of Consolidation

Suppose that you want to get started with server consolidation and virtualization. This might seem like a daunting prospect, but you can follow three simple steps to define your consolidation plan.

Identifying What to Consolidate

In a typical Exchange environment, there are a lot of moving parts aside from the Exchange servers themselves. Consider a typical Exchange 2007 environment for 10,000 mailboxes. If you follow Microsoft’s design recommendations, that environment would contain:

- two ISA 2006 servers
- two Exchange 2007 Edge Transport servers
- two (or possibly more) Windows 2003/Windows 2008 domain controllers, plus one or more AD global catalog (GC) servers
- two client access servers (CAS)
- two hub transport (HT) servers
- four mailbox servers, split into two CCR pairs, one active and one passive

That’s a substantial amount of hardware (not to mention all the associated service contracts, software licenses, and so forth). Which of these components can you effectively consolidate?

Let’s start with the infrastructure: Active Directory domain controllers and global catalogs are great consolidation and virtualization candidates. Large companies have discovered that running a 64-bit version of Windows Server allows them to have extremely large AD environments served from a much smaller number of consolidated servers, and virtualizing those servers as well (provided the design provides adequate performance) adds to the benefit.

ISA Server is also an excellent candidate for consolidation with virtualization. Why? ISA computers are essentially stateless, and they don’t store any user data. The same is true for the CAS servers, the Edge Transport servers, and even the HT servers.

Mailbox servers present a slightly different issue. There are ongoing arguments about the wisdom

Top 10 Reasons to Consolidate Exchange

10 Reduced data center space requirements. The more servers you have, the more space you have to allocate for them. The effects extend beyond just floor space, too: fewer servers means fewer racks, fewer KVM switches, fewer connections to storage, fewer connections to network, fewer cardboard boxes stacked over in the corner and reallocation of support resources to tasks more important to the business.

9 Reduced CO₂ emissions: there's ongoing debate over the degree to which climate change is human-caused, but the fewer servers you buy and run, the less you'll be contributing to CO₂ pollution, and that can't be a bad thing!

8 Positive buzz: Google has an almost uncounted number of power-hungry servers, yet they get good press for their efforts to be more environmentally friendly through consolidation and efficiency improvements. Your company can reap the same benefits (albeit probably on a smaller scale).

7 Less work: if you count up all the times you've had to drop what you're doing to attend to a faulty server, you'll probably figure out that having fewer servers will cut the amount of time you spend on emergency service calls, giving you some of your life back—and cutting your operating expenses at the same time!

6 Take better advantage of Exchange features: Exchange 2007 is designed with consolidation in mind. Its administrative tools and feature set are targeted at providing single-seat management for servers no matter where they're located, and by consolidating you can get the most out of your investment.

5 Reduced infrastructure complexity: Fewer servers deployed means an overall reduction in the complexity of your infrastructure. You'll see reduced requirements for network ports, power plugs, Active Directory sites, and other "plumbing" pieces that each require their own care and feeding.

4 Improved performance: centralizing your Exchange servers can greatly reduce the amount of RPC traffic that has to traverse your WAN, resulting in net performance improvements both for the servers themselves and the efficiency of the clients who make use of them.

3 Power and cooling savings: it's a safe bet that electricity costs will never go down in the future. Your servers convert (expensive) electricity into heat, and then you have to pay again to move that heat somewhere else. Reducing your total server count, and deploying more energy-efficient multi-core processors in those that remain, helps you save on both counts. Again, fewer is better!

2 Better operating efficiency: consolidating your Exchange workloads helps you get more utilization out of the hardware you've already bought. Having an Exchange CAS or Hub Transport server or Mailbox running at 10% CPU utilization is simply a waste at many levels of your budget.

1 Reduced capital expenditure costs: Adding a server is like buying a puppy: the initial acquisition cost is only a small part of the total lifecycle cost! Apart from the initial purchase cost, your servers need maintenance and support contracts, replacement of failed parts, and all the other expenses in the preceding nine reasons!

and efficacy of virtualizing Exchange 2007 mailbox servers. These mostly revolve around how to provide high availability, which in turn mostly involves how mailbox data is stored on disk. If you're seriously interested in consolidation, you'll almost certainly be using centralized storage as well. A virtualized set of mailbox servers combined with an iSCSI-based SAN makes a dandy combination for many applications.

Choosing a Consolidation Strategy: Physical, Virtual or Both

The basic turning point for choosing a consolidation strategy isn't necessarily the size of your environment: both small and very large environments can successfully be consolidated through virtualization. You should carefully consider several factors, including:

- How many locations you have to support. Consolidating physical servers to reduce the number of locations where you have Exchange servers deployed is a great way to begin. However, you must be careful not to design a consolidated environment where your services are put at risk by being over-concentrated.
- Your disaster recovery and business continuity needs. Exchange 2007's CCR and SCR features greatly simplify the process of setting up site failover, and you may find benefit in using a smaller number of larger, more powerful servers to consolidate and virtualize the infrastructure and mailbox servers on each "side," then use SCR and CCR to provide site failure protection.
- Your level of comfort with virtualization technology. There are still a lot of virtualization skeptics around. Like clustering before it, virtualization requires a certain amount of expertise and knowledge to manage. Until you have those, you might be better served with physical consolidation.

Consolidation and Sizing

Microsoft's position on how to size your consolidated servers couldn't be more clear: whether you use virtualization or not, you should follow the standard Exchange 2007 sizing guidelines. These are too complicated to delve into here, but they basically require you to provide an adequate number of processor cores, enough RAM, and enough disk spindles to keep disk write latency low. There's more detail in the Exchange 2007 documentation, as well as in the detailed post at <http://msexchangeteam.com/archive/2006/09/25/428994.aspx>.

What Microsoft's guidelines don't include are recommendations on the number of mailboxes you can host on a single mailbox server, nor how many virtualized servers you should pack onto a single physical host. For these, you'll need to consider the overall sizing guidelines and make sure your physical hosts and storage design are adequately sized, something that most virtualization-friendly hardware vendors are happy to help with.

An Example Environment

Let's reconsider the sample environment presented earlier. Rather than a room full of physical servers, that environment can be almost completely virtualized with a single physical host, like a Unisys ES7000, as follows:

- two VMs for the Exchange CAS role
- two VMs for the Exchange HT role
- one VM for the AD GC role
- two virtualized ISA servers
- two virtualized Edge Transport servers

Note that the mailbox server role isn't listed. That's because there are two options. One is to consolidate all 10,000 mailboxes onto a single large server, then use CCR on a second (equally large) server to provide failure protection. This is a pretty traditional path, albeit with more mailboxes per server than the norm. The other alternative is to virtualize the mailbox server role, preferably onto a separate pair of physical hosts.

Conclusion

Server consolidation has come a long way since Exchange 2003 first shipped. It's feasible to consolidate servers both by increasing the density of mailboxes (or other services) on physical servers and by virtualizing services onto larger, more powerful physical hosts. Both types of consolidation have their place, and knowing where to use them in your organization will help you gain the environmental, operational, and cost benefits that consolidation can offer.

*Paul Robichaux is a founding partner for 3Sharp LLC, an MCSE, and an Exchange MVP. He is the author of several books, including *The Exchange Server Cookbook* (Published by O'Reilly and Associates), and the creator of the <http://www.exchangefaq.org> Web site.*

<http://www.unisys.com>

multiple replicas are available, you should be able to revert to a replica database if problems occur with your live copy. Therefore, the importance of backups, especially tape backups, is lessened. Of course, administrators will have to deal with challenges such as audit requirements that might insist on offline, secured backups; providing sufficient storage and network bandwidth to handle multiple replicas and log shipping; and the inevitable updates to operational procedures necessary for backups, restores, and the loss of a disk or server.

Client Access server sizing could be another challenge. In Exchange 2007, the vast majority of Client Access workload is generated by Internet client access, including Outlook Anywhere. In an Exchange 2010 environment, the introduction of the RPC Client Access Layer means that the Client Access server has a heavier workload, so you'll find that some current Client Access configurations are undersized for the new workload.

Improved Management and Administration

Microsoft has made many improvements to Exchange's manageability, and certainly the combination of EMC and EMS in Exchange 2007 lets most administrators get their work done fast and efficiently. Both components are upgraded in Exchange 2010 to accommodate the new features and to support Windows PowerShell 2.0, which is based on Microsoft .NET Framework 3.5. PowerShell 2.0 supports remote management, so you can connect to a remote Exchange server and execute commands on it as easily as you can on a local server. In addition to new commands for features such as DAGs, some older commands are upgraded; for example, the Move-Mailbox command now supports an -online switch so that you can move mailboxes even when users are connected.

The introduction of role-based access control (RBAC) and a lightweight web console to perform a restricted set of operations are two important management changes in Exchange 2010. RBAC associates the necessary permissions with a role to let someone holding that role do his or her job effectively. We've seen the concepts of roles and associated permissions before (think of Exchange Recipient Administrator), but Exchange 2010 gives you a way to define custom roles for your organization, define the tasks that

the roles perform, and associate the permissions to allow those who hold a role to do the job. For example, you could create a Help desk role with the necessary permissions to create new mailboxes and reset passwords and such common tasks, then assign that role to the users who take care of such tasks. If you grant users the role, they automatically inherit the permissions. If the role is taken away, they lose the permissions.

Although Exchange 2007 made real improvements by optimizing Information Store caching, human behavior meant that further work was necessary for Exchange to effectively support very large mailboxes.

The big difference here is that the permissions are associated with tasks rather than AD objects such as servers and mailboxes. Thus, if you decide a role should be able to manage mailboxes, behind the scenes the role inherits the permissions required to fulfill the task. This aspect, together with the ability to set a scope of objects for a role to work with—for example, only mailboxes that belong to certain servers or only mailboxes in Germany—creates a logical and flexible approach to distributed management that should be popular with medium to large organizations. Smaller organizations will see less value in RBAC because they often have only one or two people in IT, so offloading work isn't an option.

Exchange has always included a management console, and the console includes the ability to execute tasks that are often performed by Help desk personnel, such as setting up new mailboxes or editing

mailbox properties, as well as tasks that you might not want available from the Help desk, such as creating new transport rules. Exchange 2010 adds the Exchange Control Panel (ECP), a web-based interface that lets administrators assign the ability to perform specific management tasks, using RBAC, to individuals. Smaller installations probably won't see much value in ECP, but it should be a popular feature in enterprise-class deployments.

Messaging Compliance Improvements

Microsoft created a base for messaging compliance in Exchange 2007 with messaging records management (MRM) and transport and journal rules. Unfortunately, some aspects of MRM were incomplete and difficult to deploy, such as the requirement to publish message classification definitions via XML files to each Outlook client. However, transport rules were a welcome advance, eliminating the need to write code to perform special message processing, and journal rules let Exchange efficiently capture messages. These rules depend on the architectural change Microsoft made in the transport system to force every message to flow through a Hub Transport server, even if sent to a local recipient. The Hub Transport server therefore functions as a single place where messages can be examined and processed.

Microsoft builds on MRM with some new features and by tweaking some implementation details. For example, a new records management role is defined in ECP that lets assigned individuals perform email discovery searches. Auditing will track such searches to prevent user abuse. Archiving is more granular, so you can decide to archive only messages that meet certain conditions rather than everything sent by mailboxes in a specific database or by a specific user, as is the case today. For example, you can archive messages only if the sender and recipient are in different departments or if they are located in Austria.

Exchange 2007 also introduced managed folders, each of which can have a different retention time. As it turns out, users just didn't get their heads around managed folders, so Microsoft is pursuing a different approach by focusing on tags as the basis for message retention. Administrators can

define a set of tags, such as “Important,” “Long-term archive,” or “Do not delete.” Each tag has its own retention policy, such as “Never delete these messages.” When users apply tags to messages, Exchange applies the appropriate retention policy when its management agents scan mailboxes. It’s too early to know whether tags will be any more successful than managed folders as the basis for message retention.

Exchange 2010 also includes new MRM policies so that administrators can provide

includes 3,000 recipients. Other tips will tell users when recipients can’t receive messages because their mailbox is full or if they’re out of the office and won’t be able to respond. OWA will also support MailTips and conversation views.

The Exchange 2010 Environment

Microsoft plans to release only a 64-bit version of Exchange 2010 for production, but they might again provide a 32-bit test version. Of course, now that Microsoft has

to upgrade an existing version of Exchange to the new release and will have to deploy new servers running Exchange 2010, then use the Move Mailbox feature to move users to the new servers. Details of deployment recommendations are still being worked out, but I expect that best practice will be to deploy servers running the Hub Transport (and Edge Transport) and Client Access roles first, followed by Mailbox servers.

Tons of New Developments

There are many other changes in Exchange 2010. Public folders persist, but some APIs (e.g., CDOEX, WebDAV, ExOLEDB) are replaced by Exchange Web Services. Unified messaging gains features such as a message waiting indicator and a personal auto attendant that can configure rules for how to answer incoming calls. You can expect Microsoft to connect Exchange better with Office Communications Server and its Windows Rights Management Services, bringing different strands of its information worker strategy closer together.

Microsoft still has tons of work to do before Exchange 2010 becomes a shrink-wrapped product, but all indications from the beta versions are that the new release will deliver some interesting and valuable functionality. Like any release, things can change before Microsoft ships the final software, including the elimination of features that don’t meet goals for functionality or quality. However, given that Exchange 2010 doesn’t represent the same kind of generational change represented by the move from Exchange 2003 to Exchange 2007, I expect that the bulk of the functionality that exists in today’s builds will appear in the final release. The changes in the new version collectively represent nearly three years’ hard work by a large development group, so you can expect to be busy learning all about Exchange 2010 in the coming months.



InstantDoc ID 100934

The introduction of Database Availability Groups indicates that Microsoft is heading toward multiple database replicas as the primary solution for data availability.

users with the ability to archive messages without having to move them to a PST. PSTs are horrible to deal with from an administrator’s perspective—hard to back up and restore, difficult to search thoroughly for e-discovery—so this change is a welcome one.

The Future for Exchange Clients

It’s long been standard practice for Microsoft to release a new version of Outlook alongside a new version of Exchange. Exchange 2010 is part of the Office 14 wave, so Microsoft will upgrade Outlook, Outlook Web Access, and Pocket Outlook (on Windows Mobile 7.0 clients) to add new features, improve usability, and accommodate the architectural changes in Exchange 2010, including some performance improvements within Outlook to deal with the demands of very large (>2GB) mailboxes. After all, there’s no point in Exchange being able to support very large mailboxes if its premier client finds it difficult to process those mailboxes, which is often the situation today.

The biggest thing you’ll notice in the client UI is a focus on conversation views where you’ll be able to process complete sets of messages that make up a conversation more efficiently than you can today. MailTips, small balloon-like messages, will appear to warn users whenever an action might not make sense. For example, you’re about to use *Reply to All* on a message that

Hyper-V in its armory, you can expect that Exchange 2010 will be a good candidate for virtualized deployments, albeit with the normal caveats that roles such as Client Access and Hub Transport are more suitable for virtualization than high-end Mailbox servers. Unified Messaging servers remain a poor choice for virtualization because of the demands of audio processing for voicemail. Given that experience with virtualization grows all the time, it’s wise to check with Microsoft for the latest news on support for your favorite application.

Exchange 2010 isn’t supported for Windows Server 2003, so you’ll have to deploy it on Windows Server 2008. As usual, Exchange 2010 will have other prerequisites, such as the latest version of the .NET Framework, PowerShell 2.0, and some schema updates for AD. There’s no current dependency that Exchange 2010 must access AD on Server 2008, but you’ll need to ensure that your forest is at least at Windows 2003 functional mode and that there’s at least one Global Catalog server running Windows 2003 SP2 in each domain that supports an Exchange 2010 server. Exchange 2010 doesn’t support read-only domain controllers.

Within an Exchange organization, you can mix Exchange 2010 servers with servers running Exchange 2007 SP1 or later and Exchange 2003 SP2 or later, but there’s no support for earlier versions of Exchange. Just like Exchange 2007, you won’t be able



Tony Redmond

(12knocksinna@gmail.com) is a contributing editor for *Windows IT Pro*, vice president of technology at HP Services, and author of *Microsoft Exchange Server 2007 with SP1* (Digital Press).

Scripting Utilities to **Keep Tabs** on Your Printers

These 2 scripts work together to help you keep an information history on your printers and track changes to help with troubleshooting

by **Jim Turner**

I've faced a couple of ordeals troubleshooting printer problems where I didn't have a reference to past printer configurations, so I decided to start keeping a history of all my print queues. I've found that having this history is particularly helpful in my environment because several other people have rights to maintain and troubleshoot printer problems, namely senior Help desk techs and senior desktop techs. However, if they can't resolve a printer problem, I'm the one who ultimately ends up with the support ticket. So I devised a solution for capturing my print queue data with a simple script; I can then compare the data from the current state to past states to reveal any changes and often discover problems more quickly.

A Little Background on the Problem

In addition to having eleven print servers and many printers, the number of queues in my environment is exceptionally high—more than 1,000—because many of the printers have multiple queues set up, some with PostScript drivers, some with Printer CL (PCL) drivers, and some with drivers that might appear to be incorrect because they don't match the make or model of the printer.

Having multiple techs with different levels of expertise in such an environment has its advantages. But it has its disadvantages as well. I've seen cases where someone changed settings on a print queue while trying to troubleshoot a problem, then failed to return the orig-

inal settings after discovering that the modification didn't resolve the problem. I've also frequently found that someone changed a driver from PCL to PostScript or PostScript to PCL to get a user's document to print. The tech assumes that all is well, but in fact a new problem was created for users who need the original driver.

I've seen printers that needed to have a competitor's driver installed so they would function to a certain specification that the original manufacturer's driver couldn't meet. At first sight, I can see why someone would say, "Hey, that can't be right. No wonder the user can't print." So the tech changes the driver to one that matches the printer; however, now documents don't print for the users who needed the specialized driver, and unfortunately the technician doesn't remember what the original driver was and can't reset it.

Another problem I occasionally run up against is that a print queue has been renamed, something that generally happens when



■ KEEP TABS ON PRINTERS

a printer is moved. In a facility as large as where I work, it's sometimes very difficult to coordinate migration efforts or new printer installations with all the right people, so communications occasionally get lost in the shuffle. And yet another potential problem: Sometimes I need to have information about a printer that's called back into action after being set aside in storage for a while.

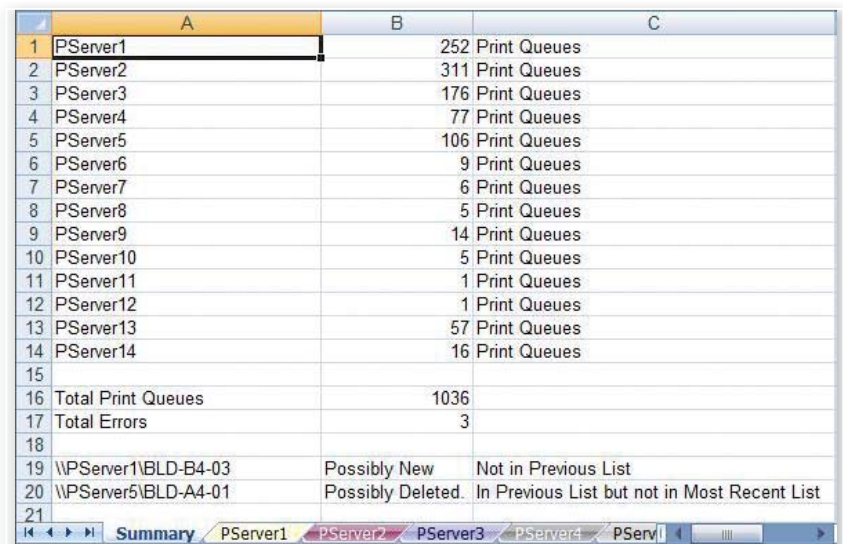
As you can see, a print queue history in my environment can be a vital resource. With as many printers and queues as we have and with possible changes being made at any time by multiple individuals, I've found that it's beneficial for me to capture printer information daily. If you work at a smaller site, you probably need to capture the data only once a week or a few times a month.

Developing the Solution

At first, I devised a solution for capturing my print queue data with a simple Windows Management Instrumentation (WMI) VBScript that gathered information from all print servers and wrote that data to a Microsoft Excel spreadsheet. This solution worked fairly well for tracking down changes; I could spot differences simply by comparing spreadsheets from different dates. But eventually I found that visually searching through numerous spreadsheets for changes, or writing macros to do comparisons, was tedious and inefficient.

Ultimately I modified my script slightly so that in addition to writing data to Excel, it wrote and saved the printer data to an ADO database as an XML file. With the data in a database, I could easily write a script that compared the data from different days in a fraction of the time it would take to do so manually.

The first of my scripts, the one I run daily, is called `PrinterInfo.vbs`; Web Listing 1 (www.windowsitpro.com, InstantDoc ID 101483) shows the code, and you can download the script from the website as well. When you run this script, it displays a spreadsheet that consists of a worksheet tab for each print server, an Error worksheet, and a Summary worksheet. The Summary sheet shows all of the print servers by name and the total number of print queues on each one. The summary also shows the number of printer errors detected on all servers combined as well as any differences between the last



	A	B	C
1	PServer1	252	Print Queues
2	PServer2	311	Print Queues
3	PServer3	176	Print Queues
4	PServer4	77	Print Queues
5	PServer5	106	Print Queues
6	PServer6	9	Print Queues
7	PServer7	6	Print Queues
8	PServer8	5	Print Queues
9	PServer9	14	Print Queues
10	PServer10	5	Print Queues
11	PServer11	1	Print Queues
12	PServer12	1	Print Queues
13	PServer13	57	Print Queues
14	PServer14	16	Print Queues
15			
16	Total Print Queues	1036	
17	Total Errors	3	
18			
19	\\PServer1\BLD-B4-03	Possibly New	Not in Previous List
20	\\PServer5\BLD-A4-01	Possibly Deleted	In Previous List but not in Most Recent List
21			

Figure 1: A sample Summary sheet from `PrinterInfo.vbs`

run of the script and the current run, such as new or deleted printers, driver changes, location or comment changes, and changes on other pertinent fields. Each of the printer worksheet tabs contains print queue information for that particular print server, and the Errors worksheet houses printer errors detected for all print servers. The first time you execute this script, of course, you won't have comparison information reported in the spreadsheet.

What `PrinterInfo.vbs` Does

As I mentioned, the `PrinterInfo` script uses WMI to gather the printer information and stores that data in an XML-based database using ADO. Here's a list of the fields that it acquires data on:

- `PrintShare` (a concatenation of `PrinterServer` and `Printer ShareName`)
- `PortName`
- `DriverName`
- `PrinterName`
- `Location` (as entered in the printer properties `Location` field)
- `DetectedErrorState` (see the `DetErr` array at callout B in Web Listing 1 for possible errors)
- `Status` (i.e., error or OK)
- `Comment` (as entered in the printer properties `Comment` field)
- `PrintProcessor`
- `PrinterStatus` (see the `PrtStatus` array at callout B for a list of statuses)
- `BiDirectionalEnabled` (i.e., true or false)
- `PrinterState` (e.g., Paper Jam, Out of

Paper; see Function `PrnState` in Web Listing 1 for a complete list)

After writing all the current data to the spreadsheet and to the database, the script opens the database from the previous run of the script (if one exists) and does an item-to-item comparison, writing any differences to the spreadsheet. First, the script compares the latest database with the previous one to find new entries; then it compares the previous to the latest to find items that might have been deleted. Finally, it compares the printers that exist in both databases to see if significant fields differ; any differences between fields are considered changed items. Figure 1 shows a sample Summary sheet.

How `PrinterInfo.vbs` Does What It Does

With the information from `PrinterInfo.vbs` readily at hand, I have a quick and accurate view of changes that took place—a history of all of my printers for a particular day and an error listing that will help me pinpoint printers that need attention. The process behind the script is relatively straightforward:

1. Set the `DBPath` variable to an existing folder that will house the databases, as the code at callout A in Web Listing 1 shows. Note that you'll need to change the path in the script to match your environment.

2. Create an array consisting of your print server names, as shown in callout B. Be sure to modify this in the script to match

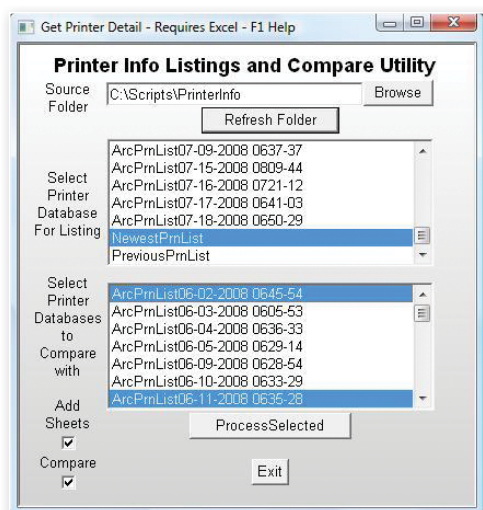


Figure 2: Comparing databases through the PrinterInfoCompare.hta GUI

your environment; simply enter all of your print servers into the array string. You'll also notice right below the print server array that I've set up arrays to accommodate the printer properties PrinterStatus and DetectedErrorState, which return only numbers; these numbers are converted to associated textual values via function calls that use these arrays before writing the data to Excel and the database.

3. Set up database filename variables, as callout C shows.

4. Create an ADO disconnected recordset with printer-related fields, as the code at callout D shows.

5. Cycle through the print servers and use WMI to collect data and write it to Excel and the database, which is what callout E shows.

6. Produce the Summary worksheet showing print queue totals and error totals, as shown at callout F.

7. Rename what was PreviousPrnList.xml to ArcPrnListmm-dd-yyyy hhmm-ss, and rename what was NewestPrnList.xml to PreviousPrnList.xml, which you can see in the code at callout G.

8. Compare databases and write any differences to the Summary worksheet, as callout H shows.

9. Save the current ADO disconnected recordset as NewestPrnList.xml, as callout I shows.

In callout C, you'll notice that in preparation for naming an archive database file, I

manipulate the DateLastModified property of the PreviousPrnList file with the functions ZeroData and MilitaryTime. This step is necessary to make accommodations for my second script, which lets you do selective database comparisons. I have to ensure that archive filenames don't exceed 31 characters, which is the limit Excel places on worksheet names. The second script, an HTML Application (HTA) called PrinterInfoCompare.hta, names worksheet tabs with the XML database filenames (minus the file extension), which makes finding specific worksheets easy. The archive filenames are in the format ArcPrnListmm-dd-yyyy hhmm-ss.xml, which satisfies the length limitation for worksheet names.

The hhmm-ss part of the filename is a military time format, which uses fewer characters, but with a hyphen substituted for what should be a colon because a colon character can't be used in a filename. The mm-dd-yyyy segment of the filename also undergoes modification. This date (as well as the time) originates from the DateLastModified timestamp of the original printer database file and doesn't usually contain leading zeros. To get the files to sort cor-

rectly for the HTA script, it's necessary to add leading zeros; so 7/7/2008 ends up as 07/07/2008, for example. You'll find the two functions, ZeroDate and MilitaryTime, near the end of Web Listing 1.

Data Listings and Comparisons with PrinterInfoCompare.hta

As time goes by, you'll undoubtedly accumulate many archive files, and there will come a time when you need to determine what changes took place between certain dates. That's where PrinterInfoCompare.hta comes into play: It provides a simple and easy-to-use GUI for performing such

- enter a path or browse to the folder where the databases reside
- select a database to get a printer data listing from
- select single or multiple databases to compare against a specific database.

You might have noticed in the title bar of the application window in Figure 2 that you can get help by pressing the F1 key. Each of the input elements on the GUI screen has its own context-sensitive Help file built into the application. You simply place the cursor into an area on the screen and press F1 to get information on that specific area. Figure 3, page 38, shows you what the Help pop-up looks like; this pop-up is also presented to users when the application is first launched. The Add Sheets check box in the GUI lets you generate multiple listings to the same Excel workbook; clearing this check box creates individual printer information and comparative listings in separate Excel workbooks.

When you select an item from the upper list box and click the ProcessSelected button, you'll be presented with a spreadsheet

There will come a time when you need to determine what changes took place. PrinterInfoCompare provides a simple and easy-to-use GUI for performing such comparisons.

showing all of the printer information available for the selected database. This output is the same as what you'd see when you ran the daily PrinterInfo.vbs script. You can produce one of these reports for as many items in the upper list box as you like. If you select an item that you've previously processed, the script opens the workbook to that particular worksheet instead of creating a new worksheet.

To compare one database with another, you must first select the Compare check box to enable the bottom list box. Select one item from the top list box, which will be your source for comparison, and one or

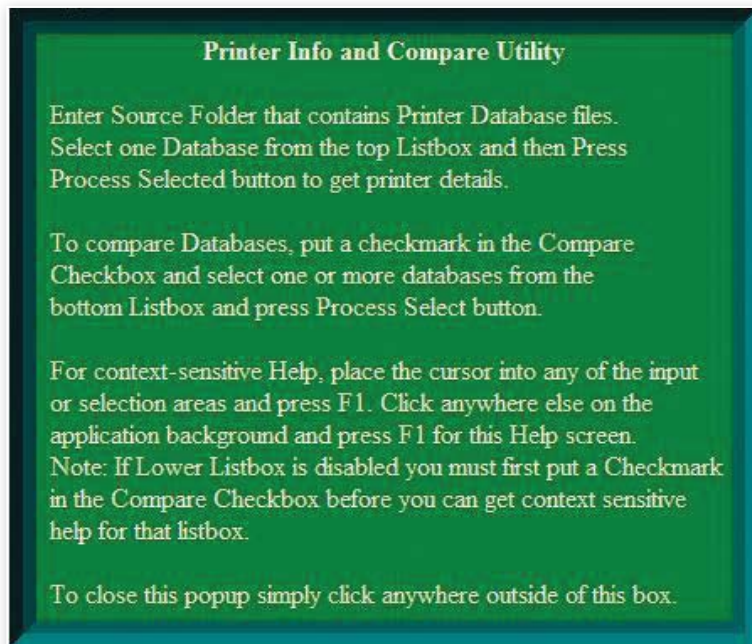


Figure 3: A sample Help screen from PrinterInfoCompare.hta

more databases from the bottom list box to compare to the source. Click the Process-Selected button to run the comparison. When the process is complete, you'll get the results in an Excel worksheet that lists any differences detected. If the databases have no differences, you'll see an entry in the spreadsheet saying "Databases Match." You'll also notice that each of the selected databases has its own populated worksheet tab so you can open that worksheet and review what that particular database has in it for printer information.

The comparison results are presented on the Compare worksheet. Information about the database selected in the upper list box is

shown in the left column of the worksheet, and information about the database you're comparing it to will be in the middle column; the last column contains driver information. When you compare the source database with multiple comparison databases, you'll need to scroll down through the Compare worksheet because comparisons are done one after the other. As Figure 4 shows, the filenames of the source and comparison databases appear above each comparison listing.

In my testing of this HTA application, I found that I could have many kinds of single listings and comparisons within the same workbook, but I don't recommend

doing that because it can become a bit overwhelming. I suggest keeping your workbooks focused on just a few comparisons. You can easily create separate workbooks within the application simply by clearing the Add Sheets check box to create a new workbook for the next listing or comparison you run. If you need to add several listings to that new workbook, just select the Add Sheets check box again before running those additional items.

Tools to Make Your Job Easier

The PrinterInfoCompare.hta script contains quite a bit of code. Rather than stepping you through all of it here, you can refer to a previous article that I wrote, "How to Easily View the Extended Properties of Files" (windowsitpro.com/article/99574/99574.html), which includes a script that performs many of the same routines as this one and contains detailed information on the code behind the process.

I certainly hope that you find these two scripting utilities useful. They should make the job of troubleshooting printer problems a little easier and help you keep track of changes and maintain printer information history.

InstantDoc ID 101483



Jim Turner

(jturnervbs@gmail.com) is a domain administrator and applications developer for Computer Sciences Corporation.

	A	B	C
1	Compare c:\scripts\printerinfo\newestprnlist.xml with:	c:\scripts\printerinfo\arcprnlist06-02-2008 0645-54.xml	
2			
3	\\PServer1\BLD-B4-03PS	Not in Compare List	
4	\\PServer3\BLD-C2-02PCL	In Compare List but not in Main List	
5	Changed	PortName	DriverName
6	\\PServer2\BLD-4-OCE4	IP_127.127.127.127	Generic 35C-1 PS
7	\\PServer2\BLD-4-OCE4	IP_127.128.127.127	Generic 35C-1 PCL
8			
9	Compare c:\scripts\printerinfo\newestprnlist.xml with:	c:\scripts\printerinfo\arcprnlist06-11-2008 0635-28.xml	
10			
11	\\PServer4\BLD-A4-01PS	Not in Compare List	
12	\\PServer3\BLD-B2-02PCL	In Compare List but not in Main List	
13	Changed	PortName	DriverName
14	\\PServer5\BLD-1-OCE2	IP_127.129.127.127	HP LaserJet 4250 PS
15	\\PServer5\BLD-1-OCE2	IP_127.129.127.127	HP LaserJet 4Si/4Si MX PS

Figure 4: PrinterInfoCompare.hta results showing comparisons to multiple databases

Create Your Own PowerShell Functions

Lesson 5 in the PowerShell 201 series explores how to create, call, and use functions

by Robert Sheldon

Like any scripting language, Windows PowerShell lets you create functions that you can reference within PowerShell statements. A function is basically a named block of code. When you call the function name, the script block within that function runs. You can include any PowerShell statements within the script block, and you can add input parameters so you can use the same function in different situations. Let's look at how to create functions, define input parameters, and work with functions in PowerShell scripts.

Creating a Function

At the most basic level, a function definition (i.e., the code that defines the function) requires the *function* keyword, the function's name, and a script block, as the following

syntax shows:

```
function <name> { <script block> }
```

The script block, which needs to be enclosed in braces, contains the statements that run when you call the function. You can include any PowerShell statement that you can run directly in the console. For example, the following code defines a function named `FileSize1`:

```
function FileSize1
{
    dir C:\Windows |
    where {$_.length -gt 100000}
}
```

Note that when you enter multiple lines of code at the command prompt, you should input a line and press Enter. You'll then see a `>>` prompt, which indicates that additional input is expected. After you've entered the entire function, press Enter twice to return to the normal command prompt (`>`).

As you can see, this function definition begins with the *function* keyword, followed by the function's name. The script block includes two commands in a single pipeline. The first command uses the `Get-ChildItem` cmdlet (represented by the *dir* alias) to retrieve the contents of the `C:\Windows` directory. The results are piped to the second command, which uses the `Where-Object` cmdlet (represented by the *where* alias) to filter out files so that only files larger

POWERSHELL FUNCTIONS

than 100,000 bytes are included in the results.

When you create a function, PowerShell stores it in memory for the duration of your session. During that session, you can call the function at any time by simply entering the function's name, as in

FileSize1

When you press Enter, PowerShell runs the code in the script block and returns the results, as shown in Figure 1. These are the same results you would receive if you had run the script block commands directly in the PowerShell console.

As this example shows, creating a basic function is a straightforward process. Although the script block here contains only a simple set of commands, you can make the script block as complex as necessary, letting you easily repeat complex logic without re-entering the same commands over and over. However, in most cases, a function without input parameters limits how much you can do with that function, so let's take a look at how to use input parameters.

Adding Input Parameters

One way you can use input parameters in a function is to take advantage of the `$args` built-in variable. When you call a function in PowerShell, you can include parameter values with the function's name. If those values aren't associated with a named parameter, they're automatically saved to the `$args` array. You can then retrieve values from that array within your function.

For example, the following function uses `$args`:

```
function FileSize2
{
```

```
    dir $args[0] |
    where {$_.Length -gt 100000}
}
```

Notice that the first command in the script block references the first value in the `$args` array (`$args[0]`) rather than specifying a pathname (e.g., `C:\Windows`). As a result, when you call the `FileSize2` function, PowerShell uses the first argument that you provide to identify the folder. If you provide more than one argument, PowerShell disregards the extra arguments because the function doesn't reference them.

To call the `FileSize2` function, you simply enter the function name and pathname, making sure there's a space between them, such as

FileSize2 C:\Windows

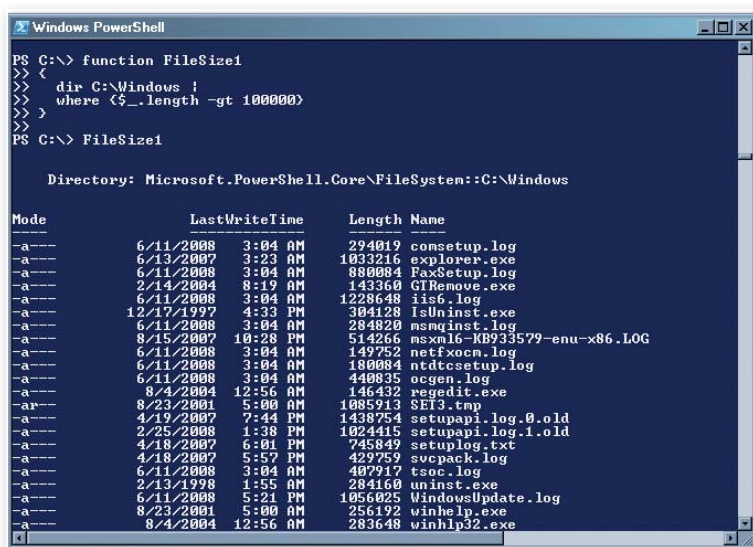
When PowerShell receives this command, it calls the function, replaces `$args[0]` with `C:\Windows`, and returns the applicable contents from that folder, providing the same results as those shown in Figure 1. Note that if a path-name includes spaces, you should enclose it in quotes.

When you call a function, each argument that you include is added to the `$args` array. As a result, you can handle any number of arguments in your function. However, working with arguments in this way can get confusing as the numbers increase. This is especially problematic if you don't enter the arguments in the correct order when you call the function. In addition, there are limitations on how you can define the arguments. As a result, creating named parameters within the function definition is often a more effective way to handle arguments.

To create named parameters, you include the parameter names, which must be preceded by dollar signs, in parentheses after the function name. When you're creating more than one named parameter, you must use a comma to separate the parameter names. For example, the function definition

```
function FileSize3 ($dir, $minSize)
{
    dir $dir |
    where {$_.Length -gt $minSize}
}
```

creates two named parameters: `$dir` and `$minSize`. The script block uses these parameters to identify the target folder and the minimum file size (in bytes), respectively.

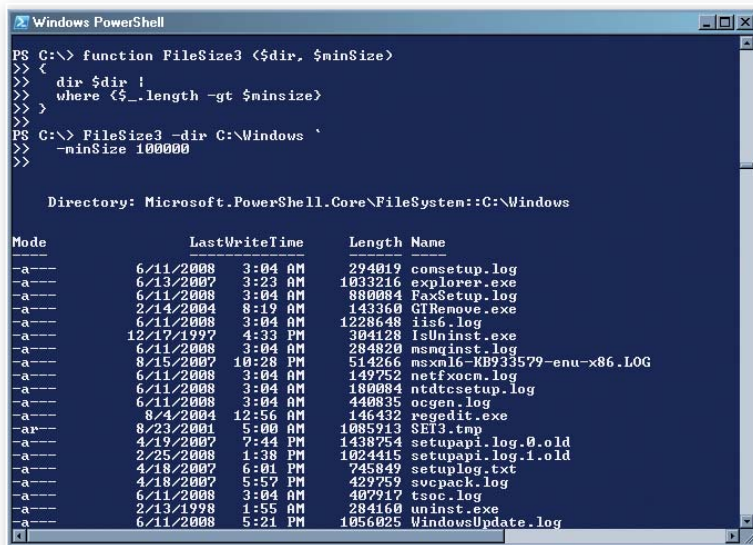


```
PS C:\> function FileSize1
> {
>     dir C:\Windows |
>     where {$_.length -gt 100000}
> }
PS C:\> FileSize1

Directory: Microsoft.PowerShell.Core\FileSystem::C:\Windows

Mode                LastWriteTime         Length Name
----                -
-a- 6/11/2008 3:04 AM      294819 comsetup.log
-a- 6/13/2007 3:23 AM      1033216 explorer.exe
-a- 6/11/2008 3:04 AM      880084 FaxSetup.log
-a- 2/14/2004 8:19 AM      143360 GTRemove.exe
-a- 6/11/2008 3:04 AM      1228648 iis6.log
-a- 12/17/1997 4:33 PM      304128 IsUninst.exe
-a- 6/11/2008 3:04 AM      204820 msqinst.log
-a- 8/15/2007 10:28 PM      514266 msxml6-KB933579-enu-x86.LOG
-a- 6/11/2008 3:04 AM      149752 netfxocm.log
-a- 6/11/2008 3:04 AM      180084 ntdcsetup.log
-a- 6/11/2008 3:04 AM      440835 ocgen.log
-a- 8/4/2004 12:56 AM      146432 regedit.exe
-a- 8/23/2001 5:00 AM      1085913 SET3.tmp
-a- 4/19/2007 7:44 PM      1438754 setupapi.log.0.old
-a- 2/25/2008 1:38 PM      1024415 setupapi.log.1.old
-a- 4/18/2007 6:01 PM       745849 setuplog.txt
-a- 4/18/2007 5:57 PM      429759 svcpack.log
-a- 6/11/2008 3:04 AM      407917 tsoc.log
-a- 2/13/1998 1:55 AM      284160 uninst.exe
-a- 6/11/2008 5:21 PM      1956025 WindowsUpdate.log
-a- 8/23/2001 5:00 AM      256192 winhelp.exe
-a- 8/4/2004 12:56 AM      283648 winhlp32.exe
```

Figure 1: Creating and running a function



```
PS C:\> function FileSize3 ($dir, $minSize)
> {
>     dir $dir |
>     where {$_.length -gt $minSize}
> }
PS C:\> FileSize3 -dir C:\Windows -minSize 100000

Directory: Microsoft.PowerShell.Core\FileSystem::C:\Windows

Mode                LastWriteTime         Length Name
----                -
-a- 6/11/2008 3:04 AM      294819 comsetup.log
-a- 6/13/2007 3:23 AM      1033216 explorer.exe
-a- 6/11/2008 3:04 AM      880084 FaxSetup.log
-a- 2/14/2004 8:19 AM      143360 GTRemove.exe
-a- 6/11/2008 3:04 AM      1228648 iis6.log
-a- 12/17/1997 4:33 PM      304128 IsUninst.exe
-a- 6/11/2008 3:04 AM      204820 msqinst.log
-a- 8/15/2007 10:28 PM      514266 msxml6-KB933579-enu-x86.LOG
-a- 6/11/2008 3:04 AM      149752 netfxocm.log
-a- 6/11/2008 3:04 AM      180084 ntdcsetup.log
-a- 6/11/2008 3:04 AM      440835 ocgen.log
-a- 8/4/2004 12:56 AM      146432 regedit.exe
-a- 8/23/2001 5:00 AM      1085913 SET3.tmp
-a- 4/19/2007 7:44 PM      1438754 setupapi.log.0.old
-a- 2/25/2008 1:38 PM      1024415 setupapi.log.1.old
-a- 4/18/2007 6:01 PM       745849 setuplog.txt
-a- 4/18/2007 5:57 PM      429759 svcpack.log
-a- 6/11/2008 3:04 AM      407917 tsoc.log
-a- 2/13/1998 1:55 AM      284160 uninst.exe
-a- 6/11/2008 5:21 PM      1956025 WindowsUpdate.log
```

Figure 2: Adding named parameters to a function

Learning Path

WINDOWS IT PRO RESOURCES

To read the previous lessons in the PowerShell 201 series, go to

"Iterating Through Collections with PowerShell's foreach Loops," InstantDoc ID 99873

"Controlling Your Code's Flow with PowerShell's Conditional Statements," InstantDoc ID 100141

"Test for Numerous Conditions with PowerShell's switch Statement," InstantDoc ID 100411

"Working with PowerShell's Data Types," InstantDoc ID 101150

When you call the `FileSize3` function, you reference the parameter's name and value the same way you'd reference cmdlet options: You specify the parameter's name—which must be preceded with a hyphen (and not a dollar sign like you did when creating them)—followed by a space and the parameter's value. If you include more than one named parameter, you simply add another space followed by the additional parameter name/value pair, as in

```
FileSize3 -dir C:\Windows `
-minSize 1000000
```

Note that, in this case, I used the back tick (`) to continue the command to a second line.

When you call a function that includes named parameters, PowerShell runs the function and replaces the parameter placeholders in the script block with the parameter values in the calling statement. For example, PowerShell replaces `$minSize` with 1000000. PowerShell then returns the result set generated by the script block, as shown in Figure 2.

If you specify the arguments in the same order as they're defined, you don't need to include the parameter names. For example, the command

```
FileSize3 C:\Windows 1000000
```

returns the same results as the command in the previous example.

Specifying Default Values for Parameters

You might find that you want the code in your function's script block to use default values if no parameter values are provided when calling that function. The easiest way to achieve this is to define the default values in

```
PS C:\> function FileSize4
>> (<$dir="C:\Windows\System32",
>> $minSize=1000000)
>> {
>> dir $dir |
>> where {$_.Length -gt $minSize}
>> }
PS C:\> FileSize4

Directory: Microsoft.PowerShell.Core\FileSystem::C:\Windows\System32

Mode                LastWriteTime         Length Name
----                -
-a-----      8/4/2004 12:56 AM      1888992 ati3duag.dll
-a-----      8/22/2007  6:12 AM      1822976 browseui.dll
-a-----      9/9/2005  6:53 PM      2067968 cdosys.dll
-a-----      7/25/2005  9:39 PM      1267200 consvcs.dll
-a-----      8/4/2004 12:56 AM      1179648 d3d8.dll
-a-----      8/4/2004 12:56 AM      1689088 d3d9.dll
-a-----      8/22/2007  6:12 AM      1054208 dani.dll
-a-----      8/23/2001  5:00 AM      1501676 diskcopy.dll
-a-----      8/4/2004 12:56 AM      1294336 dsound3d.dll
-a-----      8/4/2004 12:56 AM      1227264 dx8vb.dll
-a-----      8/4/2004 12:56 AM      1298432 dxdiag.exe
-a-----      8/4/2004 12:56 AM      2113536 dxdiagm.dll
-a-----     10/20/2005  3:20 PM      1082368 esent.dll
-a-----      8/23/2001  5:00 AM      1114896 esent97.dll
-a-----      8/23/2001  5:00 AM      1015477 esentprf.ini
-a-----     10/26/2006  2:19 PM      1190688 FM20.dll
-a-----      4/17/2007  2:32 AM      2455488 ieapfltr.dat
-a-----      4/22/2008  9:16 PM      6066176 ieframe.dll
-a-----      4/22/2008  9:16 PM      1831424 inetctl.cpl
-a-----     10/11/2007  3:12 PM      1468968 LegitCheckControl.DLL
```

Figure 3: Adding default values to parameters

```
PS C:\> FileSize4 C:\Windows 500000

Directory: Microsoft.PowerShell.Core\FileSystem::C:\Windows

Mode                LastWriteTime         Length Name
----                -
-a-----      6/13/2007  3:23 AM      1033216 explorer.exe
-a-----      6/11/2008  3:04 AM       880084 FaxSetup.log
-a-----      6/11/2008  3:04 AM      1228648 iis6.log
-a-----      8/15/2007 10:28 PM       514266 msxml6-KB933579-enu-x86.LOG
-a-----      8/23/2001  5:00 AM      1085913 SET3.tmp
-a-----      4/13/2007  7:44 PM      1438754 setupapi.log.0.old
-a-----      2/25/2008  1:38 PM      1024415 setupapi.log.1.old
-a-----      4/18/2007  6:01 PM       745849 setuplog.txt
-a-----      6/11/2008  5:21 PM      1056025 WindowsUpdate.log
```

Figure 4: Overriding default values

the function definition. For example, the following function definition provides default values for `$dir` and `$minSize`:

```
function FileSize4
($dir="C:\Windows\System32",
 $minSize=1000000)
{
dir $dir |
where {$_.Length -gt $minSize}
}
```

As you can see, all you need to do is add an equal sign followed by the default value to the parameter name. Now you can call the function without providing parameter values, as in

```
FileSize4
```

As Figure 3 shows, PowerShell automatically inserts the default values in place of the parameter placeholders in the script block.

You can easily override the default

parameter values when needed. For example, if you specify

```
FileSize4 C:\Windows 500000
```

the function returns data based on the two specified values, as Figure 4 shows.

You can also provide some values and not others when calling a function. For example, the following command includes a value for the first parameter (`$dir`) but not the second parameter (`$minSize`):

```
FileSize4 C:\Windows
```

When the function runs, it'll use `C:\Windows` for `$dir` and the default value for `$minSize`. Thus, the result set will list files larger than 1,000,000 bytes in the `C:\Windows` directory.

When you specify a parameter value in a function call that's not in the same order as the parameters defined in the function definition, you must include the parameter's

Getting Ready ^{for} Windows 7

eLearning series

with Mark Minasi

WHEN

May 28, 2009

WHERE

Your computer

COST

\$99

LESSONS

11:00 am EDT – Planning and Deploying Windows 7/Server 2008 R2

12:30 pm EDT – Operating, Managing and Securing Windows 7/Server 2008 R2

2:00 pm EDT – Directory Services Changes and Performance Enhancements

HOW

Register at www.WindowsITPro.com/go/elearning/Windows7

Is Windows 7 Right for You?

Join veteran Windows watcher and *Windows IT Pro* magazine columnist Mark Minasi in a clear, comprehensive, independent and often entertaining look at what Windows 7 can (and can't) do for you. Mark explains what's new in Windows 7 from soup to nuts to save you time and help you make an informed "upgrade or not?" decision.

INSTRUCTOR



Mark Minasi is a senior contributing editor for *Windows IT Pro*, an MCSE, and a best-selling author, popular technology columnist, commentator, keynote speaker, and IT consultant.

Mark is probably best known for his *Mastering Windows Server* and *Complete PC Upgrade and Maintenance Guide* books, both of which have seen more than 12 editions and sold over a million copies.

Learn more about the speaker, sessions, and how to reserve your seat at:
www.WindowsITPro.com/go/elearning/Windows7

Windows IT Pro

```

PS C:\> function FileSize5
>> ([string] $dir="C:\Windows",
>> [int] $minSize=1000000)
>> {
>>   dir $dir |
>>   where {$_ .Length -gt $minSize}
>> }
PS C:\> FileSize5 -minSize "file"
FileSize5 : Cannot convert value "file" to type "System.Int32". Error: "Input string was not in a correct format."
At line:1 char:19
* FileSize5 -minSize <<<< "file"
PS C:\>

```

Figure 5: Strongly typing parameters in a function

name. For example, the following command specifies only the \$minSize parameter:

```
FileSize4 -minSize 500000
```

The function will use the default value for \$dir and the 500000 value for \$minSize, so the result set will list files larger than 500,000 bytes in the C:\Windows\System32 directory. If you were to provide only the \$minSize value without the parameter name, PowerShell would assume that the value is meant for the \$dir parameter because \$dir is the first parameter defined in the function. For that reason, you must include the parameter name.

Specifying Parameter Types

In addition to assigning a default value to a parameter, you can strongly type the value by casting the variable. To do so, simply precede the parameter name with the data type name (or its alias) within brackets, as in

```

function FileSize5
([string] $dir="C:\Windows",
[int] $minSize=1000000)
{
  dir $dir |
  where {$_ .Length -gt $minSize}
}

```

```

PS C:\> $files = FileSize5 C:\Windows 500000
PS C:\> foreach ($file in $files)
>> {
>>   $file.Name + " is " +
>>   $file.Length + " bytes."
>> }
explorer.exe is 1033216 bytes.
FaxSetup.log is 880084 bytes.
iis6.log is 1228648 bytes.
msxml6-KB933579-enu-x86.LOG is 514266 bytes.
SET3.tmp is 1085913 bytes.
setupapi.log.0.old is 1438754 bytes.
setupapi.log.1.old is 1024415 bytes.
setuplog.txt is 745849 bytes.
WindowsUpdate.log is 1056025 bytes.
PS C:\>

```

Figure 6: Using a function to initiate a variable

```

PS C:\> FileSize5 C:\Windows 500000 |
>> foreach {$_ .name + " is " +
>>   $_.length + " bytes."}
>>
explorer.exe is 1033216 bytes.
FaxSetup.log is 880084 bytes.
iis6.log is 1228648 bytes.
msxml6-KB933579-enu-x86.LOG is 514266 bytes.
SET3.tmp is 1085913 bytes.
setupapi.log.0.old is 1438754 bytes.
setupapi.log.1.old is 1024415 bytes.
setuplog.txt is 745849 bytes.
WindowsUpdate.log is 1056025 bytes.
PS C:\>

```

Figure 7: Using a function in a pipeline

Now \$dir is defined with the String data type, and \$minSize is defined with the Int32 data type. If you try to enter a value with the wrong type, you'll receive an error. For example, the following command attempts to use a string as an argument for \$minSize, which is configured as an integer:

```
FileSize5 -minSize file
```

As Figure 5 shows, the command will generate an error because PowerShell cannot convert *file* to an Int32 value.

Working with Functions

Up to this point, the sample function calls have called the function directly, and the functions' results were returned to the

console. However, functions are particularly useful when used in conjunction with other elements in PowerShell scripts. For example, you can use a function to assign a value or a collection to a variable. For example, the code

```

$files = FileSize5 C:\Windows 500000
foreach ($file in $files)
{
  $file.Name + " is " +
  $file.Length + " bytes."
}

```

uses the FileSize5 function to retrieve a list of files, then assigns that list to the \$files variable. That variable is used in a foreach loop to return each file's name and size, as shown in Figure 6.

In addition to using functions to define variable values, you can use functions directly in a pipeline, along with other commands. For example, the following pipeline begins by calling the FileSize5 function:

```

FileSize5 C:\Windows 500000 |
foreach {$_ .name + " is " +
  $_.length + " bytes."}

```

The function's results are then piped to the ForEach-Object cmdlet (referenced by the *foreach* alias), which generates information about each file returned by the function, as Figure 7 shows.

Moving Forward

Functions are extremely useful when working with PowerShell scripts that perform the same tasks repeatedly. You can make your functions as simple or as complex as necessary. However, as I mentioned previously, the functions you create within a session are available only during that session. In the next lesson, I'll explain how to persist those functions so they're available whenever you need to call them.

InstantDoc ID 101610



Robert Sheldon

(contact@rhsheldon.com) is a technical consultant and author of material about Windows, relational database management systems, and business intelligence design and implementation. His latest book is *Beginning MySQL* (Wiley). Find out more at www.rhsheldon.com.

How will you manage your growing SharePoint environment?

Patch together a dozen products from a dozen vendors?

Hire a team of developers and consultants?

Get more SharePoint Admins?



Get DocAve®

One Platform for Comprehensive Backup, Disaster Recovery, Administration, Replication, Archiving, Compliance, and Migration for Microsoft® SharePoint.

DocAve by AvePoint® - Unleashing the power of SharePoint™

Find out why DocAve is the most award-winning solution of its kind
Download a FREE trial at: www.avepoint.com/unleash



Moving Your Public Folders to SharePoint

Steps to ensure a
smooth migration
by Ron Charity

During the past year, I've worked with several clients who plan to use SharePoint as a replacement for Microsoft Exchange public folders. Given the problems that organizations face with the uncontrollable growth of unstructured data, compliance requirements, and their effects on storage and operations, it's a painful topic with no "magic bullet" solution. Is there a rush to replace public folders? Probably not at the moment unless organizational factors surface that force the change. For now, there's no time pressure for organizations to switch, because public folders will continue to be supported until the end of the Microsoft Exchange Server 2007 product life cycle in 2016 or 2017. But there are compelling reasons for adopting a SharePoint-based solution: better presentation, search, and mobile access to name a few. However, expectations must be managed carefully because migrating is a complex labor-intensive undertaking that if handled incorrectly can result in significant business interruption. To help in your decision, let's talk about how to plan, design, and carry out a migration to SharePoint, with as few problems as possible.

Making the SharePoint Decision

SharePoint isn't the solution for every enterprise. Organizations that are considering migrating to SharePoint need to be aware of the following concerns:

- Moving data from public folders to SharePoint is labor intensive.
- SharePoint stores the files in Microsoft SQL Server, raising scalability concerns. Some companies have several terabytes of public folder data and file shares. For help with storage, see the TechNet article "Plan enterprise content storage" at technet.microsoft.com/en-us/library/cc263028.aspx.
- You need to tag files so that you can easily search for data. You can find guidance in the blog post "Searching Custom Column Values in MOSS 2007" at www.jjfblog.com/2007/01/searching-custom-column-values-in-moss.html.
- Moving to SharePoint can be expensive; file servers are cheaper than SharePoint and SQL Server farms.
- You'll need additional tools, such as enterprise records management and archival solutions.






Figure 1: SharePoint information mapping

- SharePoint doesn't support replication, so you'll need a third-party replication product such as a solution from Syntergy, Infonic, or other third-party providers.

Use a Fine-Grained Approach

Generally, I recommend that companies take a new approach to thinking about data and its classification, storage, and retrieval: a fine-grained approach as opposed to the "big bucket" approach of public folders and file shares. You need to consider how the information in your public folders maps to your organization's information architecture. You also need a well-thought-out destination for the public-folder data to prevent disorganization and to enable users to find the data they're looking for. After data is migrated, how will you make sure the metadata is entered for each SharePoint Content Type? For example, information contained within public folders might consist of client information, product information, job-related information, or corporate information. How do you make sure that information is transferred to the SharePoint site?

Before conducting a migration to SharePoint, you need to make sure that the public folder (and file share) migration is handled according to organizational policy to prevent compliance problems. For example, you need to determine the policy for what

data is permitted on public folders (e.g., Microsoft Office documents but not MP3s), and the policy for removal of data violating that policy. You also need to consider how the information will map between public folders and SharePoint Sites and Pages. Figure 1 provides a high-level view of how information (by functionality) maps to SharePoint. Note that it doesn't depict how information maps specific to taxonomy (detailed information architecture). In SharePoint, information is displayed in a more categorized and visible manner than it is in public folders. For example, contacts are placed in a Sites Contacts Web Part; documents are placed in a document library. Tools that migrate data from public folders to SharePoint can help with this classification and the creation of sites, but ensuring information relevance requires a lot of human involvement. Understanding this, your information architecture must address farms, Shared Service Providers (SSPs), sites, pages, Web Parts (applications), content types, metadata, labeling, security, and content organization.

I assume that you've already completed a SharePoint design that includes a detailed information architecture, system architecture, and operations plan. These pieces are crucial for decision support and for achieving a usable and compliant SharePoint navigation and search experience when you're finished. See the sidebar titled "Resources

for Planning Your SharePoint Design" if you need help with this task.

The following sections outline a basic migration methodology, including the steps for each phase and hints for dealing with your project. Keep in mind that these steps are for a very large organization with gigabytes and perhaps terabytes of public folder data. You can modify these guidelines to better suit your own organization.

PHASE 1: Project Initiation

In this phase, your goal is to build your team and prepare project documentation. These two items will help steer your project effectively.

Establish a governance team. A team and a decision framework will help you steer the rough waters ahead. Given that the project will touch just about every business unit, you'll require senior management to help facilitate the project's momentum and to address escalations and key decisions such as scope, resource scheduling, data retirement, prioritization, and business unit buy-in, to name a few. To be effective, the governance team should consist of executives from IT and business units, IT architects, the project management office, and purchasing.

Develop a project charter. As with any project that deals with information, scope creep is your enemy because it can increase complexity and result in lengthy project schedules. Agreeing upon scope and priority is difficult especially if your requirements cover migration, security, and data cleanup. Your communication plan must address who, what, and when, and the business units that are to be migrated should be communicated with early in the process. Expect push-back due to concerns about business interruptions.

PHASE 2: Requirements

In the requirements phase, you develop a

document and supporting materials that define the specific requirements of your organization and lead your team through the process of developing a site design. For example, your requirements might consist of migrating public folder data to SharePoint sites, establishing security guidelines, configuring SharePoint security accordingly, and ensuring compliance with information-management policies and information architecture. The requirements document should contain the requirements for the team, tools, methodology, and risk plan. When developing requirements, remember to include the following components.

Project Management Office (PMO).

The PMO will have some insight to the project's initial scope, deliverables, and time lines.

Compliance department. The compliance department might already be actively involved because of past audits that exposed compliance issues, but if not, I suggest you meet with them to develop a list of compliance requirements. This could be as simple as a list of principles that must be incorporated into a design (e.g., being able to identify documents that reside on public folders that have legal impact, such as contracts).

IT department. You need to consider IT requirements for infrastructure and operations, such as capacity requirements, and take into consideration Help desk, monitoring, and n-level support. Note that it's important to involve the Help desk people since they'll deal with support calls after the data migration begins.

Business requirements. Though often avoided by IT departments, working closely with the business early in the process is critical to managing perceptions and determining their specific needs. Also, don't use the business as your testing ground. Use a lab and build a mockup of a business unit for testing.

Quality Assurance (QA). Many organizations have a QA process that can add significant time to your document acceptance process. Don't forget to factor this into your timeline and documentation plan. Meet with the QA people to understand what you must provide them and when.

Inventory. The first major task for you will be to develop an itemized inventory of

Resources for Planning Your SharePoint Design

Before you can migrate to SharePoint, you need to create a SharePoint design that includes detailed information architecture, system architecture, and an operations plan. These pieces are crucial for decision support and for having a usable and compliant SharePoint navigation and search experience when you're finished. If you need help planning your SharePoint design, the following resources are available:

- Information Architecture: See "Logical architecture components," technet.microsoft.com/en-us/library/cc263121.aspx, and "Information architecture in Office SharePoint Server," technet.microsoft.com/en-ca/library/cc262985.aspx.
- System Architecture: See "Planning and architecture for Office SharePoint Server 2007," [technet.microsoft.com/en-ca/library/cc261834\(TechNet.10\).aspx](http://technet.microsoft.com/en-ca/library/cc261834(TechNet.10).aspx); "Planning for Capacity Boundaries, Estimating Performance & Capacity Requirements... Additional Factors, and Tools," blogs.msdn.com/joelo/archive/2006/11/22/planning-for-capacity-boundaries-estimating-performance-capacity-requirements-additional-factors-and-tools.aspx; and various HP white papers at <http://h71019.www7.hp.com/ActiveAnswers/cache/70675-0-0-0-121.html>.
- Operations and Governance: See "Governance Resource Center for SharePoint Server 2007," technet.microsoft.com/en-us/office/sharepointserver/bb507202.aspx.

the public folders. To be successful, you'll require a toolset that can crawl and inventory the public folders and provide robust and customizable reporting. The tools you choose must be installed early on in the project so that inventory of the public folders can begin. Several vendors make tools for migrating public folders to SharePoint, including Quest Software, Metalogix, and Tzunami.

Analysis. Generally your analysis will focus on the following: What data do you have and how much? Where are the security and compliance risks? What data can be deleted to reduce storage and operations costs? What data can be reused and placed in SharePoint?

Other projects. Most organizations have several projects underway at the same time. You must plan for this because collisions will occur and dependencies must be understood. Meeting with the PMO will help you understand what projects are planned or underway.

PHASE 3: Design

During the design phase, you develop a document and supporting material to define the specific elements of your design. The design document contains (depend-

ing on your organization's methodology) the approach, methodology, and support materials for your migration. This document should have tight linkages to your project requirements document and must address how the requirements listed in the requirements document will be addressed. The document should include the following items.

Tools infrastructure. The toolset you use to inventory the public folders ideally has the ability to migrate the contents to SharePoint sites and log the results. If not, you must assess toolsets based on the requirements document. Your design must incorporate the technical infrastructure required to support the toolsets for the duration of the project. For example, how many servers do you require for the toolset? Do you need a workstation to act as the operator console? Does the toolset require a database? How much storage does the database require? Are agents required on the servers? What are the impacts to network bandwidth? For example, one organization's network between the United States and Australia didn't support the bandwidth required to migrate data. To address this issue, the server's drives were removed and shipped to the United States and the

■ PUBLIC FOLDERS TO SHAREPOINT

migration was performed there. If time permitted, perhaps a replication approach might have been a better solution.

Application remediation. Your design must incorporate tools and processes for dealing with applications that rely on the public folder infrastructure. Generally, the tools should be able to spot such applications; otherwise you must have a manual inventory process for collecting such information from the business units. There must also be a central record (e.g., a spreadsheet or database application) of applications to be remediated. Last but not least, your design should include a design for a development and QA environment for recoding and testing rewritten applications.

Public folder to SharePoint mapping. How will public folder data map to your SharePoint environment? Create a form in Microsoft Excel that lists the mapping of the folders to SharePoint and any specific notes such as exclusions. Note that migration tools will attempt to create sites based on the public folder hierarchy and populate those sites with the content types contained within them. Expect to do some cleanup once the migration process is completed.

Work package. A work package contains a summary of the work assignment and any forms or checklists the user will require while performing the work.

Staffing. This section should describe the staffing model required to execute the public folder migration. Also, skill sets and experience must be listed here.

Training. The document should include training requirements for Help desk staff, IT staff, and SharePoint end users.

Testing. A well-defined test plan with scenarios and "How to..." checklists is a must to ensure that the migration occurred as planned. For example, you should perform the following tests:

- Data migration: Determine whether the public folder data was migrated and the target SharePoint site is in place.
- Security: Make sure the desired security model is in place.
- Data policy: Ensure that the data that does not conform to policy hasn't been migrated.
- Search/browse: Make sure that the data can be browsed or searched using SharePoint.

Risk management. A common way to address risk is to develop a risk plan document that lists each of the risks. An approach that works well is for you and your team to identify the risks, then analyze and rate each according to probability (high, medium, or low) and impact (high, medium, or low).

When you're developing your design document, lean on your team on a regular basis for design advice, reviews, and sanity checks. Ultimately you should have a weekly discussion with them so that you can stay informed about each others' projects, tasks, and roadblocks. You also want them to be coauthors for the document so they have some skin in the game.

PHASE 4: Application Remediation

This phase deals with the remediation of the applications discovered through the inventory process conducted with the business units during the previous phase.

Is there a rush to replace public folders? Probably not at the moment. But there are compelling reasons for adopting a SharePoint-based solution.

Generally, this is the most time-consuming process; each application is assessed to determine its specific requirements such as technology and level of effort. Generally, applications can be categorized as low, medium, or high complexity according to the following guidelines:

- Low: Your toolset and process have identified a simple solution. An example of low complexity is simple code

changes that remove specific public folder-related APIs and replace them with SharePoint-related APIs.

- Medium: Your toolset and process have identified a solution that involves moderate recoding and testing. An example of medium complexity is simple code changes that remove specific public folder- and third-party product or Line of Business (LOB) applications-related APIs and replace them with SharePoint-related APIs.
- High: Your toolset and process cannot identify a solution and therefore more detailed assessment is required. For example, the application requires recoding and additional products such as Office InfoPath to provide forms and SQL Server for a data repository.

I highly recommend that you perform application remediation early in the project. In large organizations, this phase should occur perhaps six to eight months in advance of the data migration phase.

PHASE 5: Pre-Migration

This phase is all about making sure you (and your users) are ready to undertake the public folder migration. This phase is mostly about managing quality and risk. Also note that the pre-migration steps are specific to the toolset you choose. (For example, reporting is automated or a manual process.) Before migrating, check that you've completed these steps:

1. Communicate your plan to IT and the business.
2. Update your requirements and design documentation to reflect the realities of your organization.
3. Add resources or make project team changes based on how well people are working together, workload, and scope.
4. Schedule migration jobs to run at predefined times (off hours); use caution to schedule jobs outside of other resource-intensive jobs such as backup, virus scanning, and indexing. Tools such as Quest Migrator offer flexible job scheduling.
5. Give IT an onsite presence to provide support, especially for complex requirements and high-visibility business users. It's surprising how many companies forgo this.

SharePoint Showcase

Your one-stop resource for SharePoint products and services



DocAve v5—The world's most powerful and award-winning solution for SharePoint backup, disaster recovery, administration, replication, archiving, compliance, and migration.

Unleash the power of SharePoint!

www.avepoint.com



Increase profitability with the world's largest SharePoint Hosting Provider. Rapid Deployment. SAS70, SOX, HIPAA. Trusted by US Air Force.

The SharePoint Hosting Pioneers

www.fpweb.net/sharepoint-hosting



Cost-Effective Solutions to Accelerate SharePoint Adoption

Realize the benefits of SharePoint in this web seminar, which will explore migration approaches and best practices to more effectively transfer content from Lotus Notes and QuickPlace onto SharePoint 2007.



Register for this **FREE** resource today at www.windowsitpro.com/go/WebSeminar/SharePointAdoption

Best practices for Deploying Microsoft Office SharePoint Server 2007



In this white paper, learn best practices along with considerations for deploying Office SharePoint Server 2007 with Hyper-V on HP ProLiant servers and HP BladeSystem servers with sample configurations for typical Office SharePoint Server 2007 deployments when using Hyper-V.

Register for this **FREE** resource today at www.windowsitpro.com/go/WhitePaper/SharePointDeployment

AvePoint Web Seminar

Comply with Complete Confidence

Protect your organization's SharePoint deployment from the risk of unauthorized use and non-compliance with regard to policies and regulations. In this web seminar, you'll learn about compliance-related risk factors associated with unsuccessful SharePoint deployments, enforcement of retention and compliance policies via archiving, and how to monitor SharePoint access and usage.

Register for this **FREE** web seminar today at windowsitpro.com/go/ComplywithConfidence

Get Full Access That Fits Your Schedule

with the
WindowsITPro

Monthly Online Pass

For only \$5.95 per month, you'll get access to over 10,000 Windows IT Pro online articles, including Dan Holme's expert SharePoint solutions! Plus, you'll get the latest digital issue of Windows IT Pro magazine!

Sign up today to start getting the answers you need when you need them.

www.windowsitpro.com/go/MonthlyPass

Ease SharePoint Migration and Ongoing Site Maintenance

Learn about the various migration options you have and a valuable tool that can help with both migration and the ongoing site maintenance often required by content owners and SharePoint site admins—all in this white paper written by independent SharePoint expert Michael Noel.

Simplify your migration task—download this FREE white paper today at

WindowsITPro.com/go/EaseSharePointMigration

6. Create a master record of applications to be remediated. This list is co-developed between IT and the business and includes rewrites and applications to be replaced by a consumer off-the-shelf (COTS) product. The list must contain contact information and some sort of complexity rating.

7. Train IT staff in SharePoint installation and administration, and make sure that business units also are trained in general how-to and company-usage policies prior to the actual migration.

PHASE 6: Migration

This phase consists of the steps in the actual migration of public folder data to SharePoint sites. The actual steps for conducting the migration will depend on the migration toolset you've chosen because screens and options will be different. Therefore, I'll just generalize the basic steps, as follows:

1. Notify IT that the migration is about to occur and notify the business unit that their data is about to be migrated.
2. Establish onsite presence and ready the Help desk.
3. Using the migration job schedule you created during the design phase, create the migration jobs, and schedule them to run accordingly. To support your test plan, make sure you enable logging so that when the migration is complete, you can check for errors and deal with them. Also, use caution when configuring the security aspect of the toolset; going with minimum permissions is probably the best approach. Finally, set filters to prevent the migration of data that doesn't conform to your organization's policy.
4. As the jobs run, monitor the jobs and the performance of the servers, storage, and the network. As the public folder data is migrated, it will tax these systems significantly unless the toolset provides throttling settings.
5. Execute the test plan you created during the design phase.
6. During and after migration, communication must be rigorously maintained between the migration team and the Help desk. Debriefing with the Help desk after migrations are complete for a business unit will help you learn and refine your methodology. Review Help

desk incidents to learn where improvements could be made. Also, expect some cleanup work to be done to fine tune the organization of sites and data. Depending on your organization's information architecture and expectations, this could be a lengthy process.

Note that when escalation is required, you will require a process and clear ownership of tasks. From a governance perspective, you'll need a process for engaging with management in case you require their guidance or authority to obtain a decision or facilitate an action. For an excellent book to help you plan your governance program, see Peter Weil and Jeanne Ross's *IT Governance: How Top Performers Manage IT Decision Rights for Superior Results* (Harvard Business School Press, June 2004).

PHASE 7: Post Migration

During the post-migration phase, the organization is charged with maintaining the information architecture and enforcing the information management policy. Here are steps your organization can take to facilitate the success of these tasks:

1. Establish monitoring and reporting processes and tools to ensure data quality, information architecture compliance, and security compliance. Most tools have predefined reports to help with reporting; the time-consuming aspect of this task is reviewing reports and escalating issues to management.
2. Assuming SharePoint is new to your organization, you will have to ramp up staff and outfit your IT infrastructure with backup, monitoring, virus scanning, and other tools. And don't forget about the SQL Server team: Farm databases require regular maintenance to maintain performance. See the Microsoft article "Database Maintenance for Microsoft SharePoint Products and Technologies," which provides valuable information for SQL maintenance specific to SharePoint (office.microsoft.com/download/afile.aspx?AssetID=AM102632301033).
3. Educate IT about the changes in technology and how they affect the services they provide. Impacts to SLAs and operations must be communicated and understood.

Learning Path

WINDOWS IT PRO RESOURCES

"How SharePoint Matches up to Public Folders," InstantDoc ID 96139

"Managing Public Folders in Microsoft Exchange Server 2007," InstantDoc ID 97145

"SharePoint Server 2007 Unleashed," InstantDoc ID 94652

"Strategies for Migrating Public Folders to SharePoint," InstantDoc ID 96744

4. Educate staff and management about the changes in application technology and how they affect their jobs and their responsibilities.

Involve Human Resources to make sure that both IT and staff use the educational training. Also, to facilitate user adoption, usage metrics should be added to the job descriptions of users so that usage in compliance with company policy can be measured. For example, project managers are responsible for uploading project-related artifacts such as charters, schedules, and design documents.

Making the Leap

Microsoft's investment in public folders has noticeably declined in recent years, and SharePoint is clearly being positioned by Microsoft as the replacement platform. SharePoint offers many comparable features in addition to providing a platform for building, deploying, and managing applications. The decision to migrate your public folders to a platform such as SharePoint depends on your organization's information strategy and such factors as the complexity of the current deployment and the availability of the necessary funds and resources.



InstantDoc ID 101412



Ron Charity

(Ron.charity@hp.com) is an HP solution architect with 20 years of experience. He's located in Toronto, Canada, where he focuses on solutions for document and records management, collaboration, search, portals, and social networking.

PRODUCT SPOTLIGHT

Citrix Explains the New Essentials

Citrix announced the details today of how it will do business now that XenServer is a free product: the company will sell a new product, **Citrix Essentials for Hyper-V and XenServer**. The free version of XenServer isn't a crippled version. It includes all the features that were included in XenServer Enterprise other than advanced high-availability functions. Simon Crosby, CTO of Citrix's virtualization and management division said that unlike its competition, the free version of XenServer provides much more than a hypervisor.

Essentials will extend the abilities of both Hyper-V and XenServer—Crosby said that though there is some occasional overlap between the companies' products, Citrix is focused on extending the capabilities of Microsoft products instead of competing with Microsoft. The two hypervisors are very compatible—a server can be moved from XenServer to Hyper-V, though not as a live migration.

Essentials adds high-availability functions but also adds new features, such as StorageLink, which manages storage so that virtual machines (VMs) can have full access. Crosby said storage management is what people will need from virtualization products now that workloads can be moved around easily. "I think a hypervisor is easier," Crosby said. "Storage is hard."

Essentials will also support dynamic provisioning, allowing you to boot multiple systems from a single image. Essentials' image manager will be able to dispense machine images not only to be run as VMs but also to boot client and server system hardware. The Platinum Edition of Essentials will also add automated lab management for easier testing in virtual environments. To learn more, call 954-267-3000 or visit www.citrix.com.

■ Virtualization ■ Exchange

Exchange 2007 Backup on Server 2008 for SMBs

Cortex I.T. has released **BackupAssist 5.2**, which includes a plug-in to provide backup of Microsoft Exchange Server 2007 on Windows Server 2008. You can schedule BackupAssist to perform a variety of pre-defined backups or customize the solution to meet your environment's needs. BackupAssist includes a full array of reporting and notification options—also customizable—such as email, network broadcasts, and printouts. You can even have it remind you when it's time to clean the heads on your backup tape drive. BackupAssist works with a variety of backup mediums—tape, external drives, CD or DVD, and others. The product is priced beginning at \$249 and the Exchange Mailbox Add-on is \$129. For more information, visit www.backupassist.com.

Safari 4.0 Beta In Action

Apple has released **Safari 4.0**. The new browser matches Google's Chrome in a lot of ways—the tabs are at the top of the window, and Safari opens new tabs with a 12-panel view of your most visited sites. Like Chrome, Safari's default appearance is very sparse, and it also has a favorite sites page (see the image). Page loads were very

■ Security



fast. Unfortunately, some sites, such as Hotmail and Microsoft sites, will not behave properly in Safari. To learn more, visit www.apple.com/safari.

I'm InTouch SecurePC Prevents Lost Data on Stolen Laptops

Remote Access Software has released the **I'm InTouch** secure access platform. This solution includes the I'm InTouch SecurePC remote access terminal, SecureKEY physical authentication USB key, and I'm InTouch remote access service. The I'm InTouch remote access platform safeguards sensitive data behind the corporate firewall; remote users access

their workstations through the SecurePC remote access terminal and the I'm InTouch remote access software. One SecurePC with a SecureKEY and a year of remote access is \$700. For more information, call 800-668-2185 or visit www.01com.com/secure.



Jeff James | jjames@windowsitpro.com

Editor's Note: Send new product announcements to products@windowsitpro.com.



Providing the **vision** and **intelligence**
to keep you and your company
competitive in today's market!

One **Place**, One **Time** ...

WINDOWS
CONNECTIONS

MICROSOFT
EXCHANGE
CONNECTIONS

UNIFIED
COMMUNICATIONS
CONNECTIONS

SharePoint
CONNECTIONS

SQL SERVER
CONNECTIONS

Mobile
Mobile PC + Mobile Devices
CONNECTIONS

MICROSOFT
ASP.NET
CONNECTIONS

VISUAL STUDIO & .NET
CONNECTIONS

WinConnections Fall '09

NOVEMBER 9-12, 2009

LAS VEGAS, NV • Mandalay Bay Resort & Casino

WINCONNECTIONS ROCKS the technology industry as the **LARGEST** and most **EXCITING** event focused on **MICROSOFT TECHNOLOGY** and **YOUR** needs.

Exciting Announcements:

Be among the first to get the insiders scoop on the products and technology you rely on!

As a WinConnections attendee, you and your colleagues can attend all of the Connections shows, and cross between all of the sessions, at the same time for the same price.



Steve Riley
Microsoft
Senior Security Strategist
in Microsoft's Trustworthy Computing Group



Mark Minasi
MR&D
Best-selling Author,
Popular Technology Columnist,
Commentator



Scott Guthrie
Microsoft
Corporate Vice President, .NET Developer Division



Thomas Rizzo
Microsoft
Director, SharePoint Group

>> The first 400 paid attendees will be mailed SQL Server 2008 standard with one CAL

CHECK WEB SITE FOR DESCRIPTIONS OF SESSIONS AND WORKSHOPS

www.WinConnections.com • 800.505.1201 • 203.268.3204 • Register Today!

Microsoft®

TechNet

TECH
Conferences
PENTON MEDIA

WindowsITPro

Xobni

Read the full-length review at www.windowssitpro.com, InstantDoc ID 101726.

When I read Siegfried Jagott's article about **Xobni** (see the web-exclusive article "New Add-On Changes the Way You Work with Outlook," May 2008, InstantDoc ID 99326), the author's rave review of the tool and its comprehensive mailbox indexing and search capabilities inspired me to try it out. I went to the Xobni website, entered my email address, clicked the Download button, and followed the instructions to download and install the free Xobni tool on my Windows Vista laptop. (Xobni runs on Windows XP SP2 and later and Outlook 2003 and later.) The entire process took about 15 minutes, including restarting Outlook after the Xobni setup.

The first thing I did was resize parts of my Outlook UI to fit in a minimized Xobni pane. Then I clicked the expand button at the bottom of the pane, to open up the Xobni initial view, which Web Figure 1 (www.windowssitpro.com, InstantDoc ID 101726) shows. Clicking a message in my Outlook

Inbox displayed a pane containing messages, names, and attached files associated with the message's sender.

Xobni helps you quickly find specific messages associated with your search keyword—a name or topic, for example. Although Outlook also lets you search by name or subject, Outlook's search results don't immediately reveal what part of the message contains the search keyword. Xobni does a better job of displaying the relevant parts of messages containing the highlighted search keywords you enter.

Next, I tried out the search bar—a feature I use frequently in Outlook. I wanted to quickly find the latest version of *Windows IT Pro's* 2009 editorial calendar, so I typed "editorial calendar" in the search bar. Xobni starts searching as soon as you type in a complete word; I didn't need to press Enter. It took about 1 second to display all results in my All Mail Items folder, compared with 4 seconds using Outlook's native search. Xobni also didn't require me to navigate to All Mail Items. Clicking the message

displayed a Xobni view of the message, including the attachment. If I had wanted to open, reply to, or forward the original Outlook message, I could have done so by clicking any of the links at the top of the message, which Web Figure 2 shows.

Xobni complements and enhances Outlook and integrates well with the Outlook UI. My only complaint is that Xobni doesn't search or integrate with Outlook Contacts. If you want to get the most out of Outlook, especially its search facility, I strongly recommend adding Xobni.

InstantDoc ID 101726

Xobni

PROS: Easy to install and use; fast, accurate searching

CONS: Doesn't include Outlook Contacts in search

RATING: ◆◆◆◆◆

PRICE: Free

RECOMMENDATION: If you already like Outlook, Xobni will help you love it.

CONTACT: Xobni • 415-986-5101 • www.xobni.com

PatchSee

Read the full-length review at www.windowssitpro.com, InstantDoc ID 101641.

The **PatchSee** system, sold in the U.S. through Mitsubishi International Corporation, aims to clean up the tangle of Ethernet cables in your server room. PatchSee cables feature optical fiber that runs inside the cable, allowing you to identify them without disconnecting them. You shine light into one end of the cable using a special tool and the other end lights up.

The lights identifying the cables are small but clearly visible in most situations. It was easy to find the other end of a cable in a well-lit room, but direct sunlight overpowered the light from the cables. In Figure 1, the blue dot on the cable in the 3X slot is the light from the other end of that cable.

Also visible in Figure 1 is the cables' other distinguishing characteristic, PatchClips. These colored plastic clips snap firmly onto cable jacks. I don't see the point of PatchClips—they don't do anything you



Figure 1: Four PatchSee cables

couldn't do with colored tape.

An important limitation to PatchSee cables is that you can't cut cables to the lengths you need as you can with standard Ethernet cables. PatchSee cables also cost substantially more than standard cables, and you can't buy them in bulk lengths.

If you can tell where all the cables connect in your server room, PatchSee cables don't offer much for you. If you frequently

attach and remove Ethernet cables or your server room looks like a mass of spaghetti, however, these cables could be invaluable.

InstantDoc ID 101641

PatchSee

PROS: Allows tracing cables without risking disconnection; simple to use; light from optical fibers is very visible

CONS: Cables are expensive compared with standard cables and limited to fixed lengths; PatchClips are hard to see from some angles

RATING: ◆◆◆◆◆

PRICE: PatchSee system starts at \$45; 5' Cat 5e cables cost about \$6; 50' Cat 6 cables cost about \$30

RECOMMENDATION: If you already have a system for keeping your Ethernet cables organized, you don't need PatchSee, but if you frequently need to figure out which cable leads where, PatchSee could be a lifesaver.

CONTACT: PatchSee • 44-0-208-777-6161 • www.patchsee.com



Anne Grubb | agrubb@windowssitpro.com
Zac Wiggy | zwiggy@windowssitpro.com

Unbounded Printing Services for SharePoint

Recently, SharePoint was deployed at your company. Soon after, the Help desk receives a call from the CEO's assistant, asking, "How do I print the different versions of a document without restoring and opening each version?"

The Help desk tech says, "I'm not sure. I'll have to look that up."

Then Harry from sales calls and asks, "Do I have to open each document in order to print it?"

Once again, the Help desk tech isn't sure and says, "Well... I've never tried that. Let me take a look and call you back."

After looking through the SharePoint Help menus and doing a few Google searches, it becomes clear to the Help desk tech that these functions aren't part of the standard SharePoint deployment. So then the tech has to call the users back and tell them "You can't do that with SharePoint."

SharePoint utilizes network and workstation print resources. Therefore, when you open files from a document library, you can print only one document at a time. Often, there are more expanded printing needs with documents. **Unbounded Printing Services for SharePoint**, a product from Unbounded Solutions, enhances SharePoint's printing abilities by integrating print features into the SharePoint site itself. (Unbounded Printing Services for SharePoint is available in three versions: Small Business Edition, Standard Edition, and Enterprise Edition.) The product doesn't create a print service that supersedes the Print Spooler services, but rather adds the printing abilities that are often requested by users.

Print Services

I found the installation of Printing Services for SharePoint to be straightforward. This product requires Windows SharePoint Services 3.0 or Office SharePoint Server 2007, Windows Server 2003 or later, Internet Explorer 6.0 or later, and the .NET Framework 2.0 or later. In addition, end users must have Microsoft Office 2007 software locally installed on their workstations to process the file types. The product's features aren't deployed throughout the SharePoint farm by default, so you must enable the designated print services

on each site, which gives you control over the printing abilities within each site.

I realized while working with the product that a good deal of thought was given to the structure of the print services. Printing Services for SharePoint provides both granular and bulk control over data at the list level, folder level, and item level. With regard to document properties, Printing Services for SharePoint lets users print out the properties of a file from within the document library (as shown in Figure 1) and the version history of one document, or even all the documents, in the library. This feature lets users track the changes to key business documents, such as in a biotech project where complex processes are documented, with contributions being made to the document by a project team.

Often a SharePoint list will receive email messages with attachments. Printing Services for SharePoint lets the SharePoint list users print some, or all, of the attachments without having to open each one. This product also lets users print single or multiple InfoPath forms without having to open the forms one at a time.

With Printing Services for SharePoint, users can print only documents and versions that they have access to, so site security isn't compromised by the power of the printing features. In addition, users can create snapshots of SharePoint libraries and lists. It also offers enhanced calendar printing that lets users print event properties and views without manually opening each event.

Because these print services let you more easily print and offer quicker access to items, it seems clear that this product can contribute to a quick adoption of SharePoint by users throughout the enterprise. Some things should be considered before implementing this product in your environment, though. Printing Services for SharePoint doesn't supersede your network's print services; it still depends on them for printing.

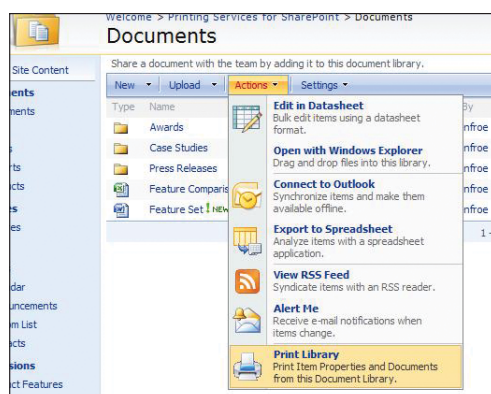


Figure 1: Printing item properties and documents

So if there's a problem with your organization's base printing resources, Printing Services for SharePoint won't replace or improve them. This product also doesn't speed up printing. So you need to take your print infrastructure and the bandwidth of your network into consideration. This is particularly true if end users are planning to bulk print SharePoint items from a remote office.

Increase Productivity

When I reviewed this product, I was immediately reminded of how my enthusiasm for SharePoint was dulled by its printing limitations. If your organization intends to make extensive use of InfoPath forms, calendar events, list item attachments, and reports, then I suggest giving Printing Services for SharePoint a try. SharePoint has the potential to increase business productivity if it isn't taxed by an insufficient printing environment.

InstantDoc ID 101649

Unbounded Printing Services for SharePoint

PROS: Great way to meet specific printing needs in SharePoint; increases productivity

CONS: Relies on existing print infrastructure

RATING: ◆◆◆◆◆

PRICE: Starts at \$2,275 for Small Business Edition

RECOMMENDATION: I recommend this product for midsized to enterprise SharePoint deployments.

CONTACT: Unbounded Solutions • www.unboundedsolutions.com • 412-571-6377



Curt Spanburgh | osgcurt@onesolutiongrp.com

KVM over IP Switches

Control your infrastructure from anywhere you want

by Jason Bovberg

[Editor's Note: To view this month's buyer's guide table, visit www.windowsitpro.com and enter 101689 in the InstantDoc ID text box at the top of the page.]

The KVM over IP switch is one of the most fundamental components in your enterprise network infrastructure, and it's clear how the functionality can improve your efficiency: It gives you in-band or out-of-band access to system keyboard, video, and mouse (KVM) functions from any location at any time. In our July 2007 issue, we presented a KVM over IP switches buyer's guide (InstantDoc ID 96095) that showcased products of major vendors in the field. We've revisited the market by sharing some new offerings from favorite vendors and introducing you to some newcomers.

What KVM over IP Brings You

Particularly if you head up a sprawling IT environment, you face the challenge of overcoming geographic barriers in your day-to-day network management; you need to react to problems on far-reaching systems as quickly as possible. Or if—in the clutches of our economy—you're performing solo IT administration, you need to increase productivity despite your lack of resources. A KVM over IP switch lets you maintain and manage geographically diverse devices, better manage systems to deliver key business services, and reduce total cost of ownership. KVM over IP switches give you BIOS-level control of connected servers and other network devices straight from any location: From a central interface, you can securely manage your entire IT infrastructure—including branches and remote data centers—as if you were administering them locally. A good KVM switch gives you complete access to authentication, event alerts, and user log files. Some KVM solutions even let you manage all your servers and devices when the network has failed and remote-access software isn't functioning.

Purchase Factors

KVM over IP switches can differ substantially in their breadth of functionality. To avoid wasting valuable resources or even compromising your business's security, you need to consider carefully the options you need for your unique environment. For example, the solution you choose needs to be able to support every OS platform and network device in your environment. Most of the solutions in Web Table 1 support a broad range of platforms. You might not have

some of these platforms in your local environment, but don't forget that your network probably knows no boundaries; you must also consider remote users' laptops and mobile devices.

How many ports do you want the switch to have? As your company inevitably grows after this downturn, you'll need it to handle more than it needs to handle now. Switches differ widely in the number of computers that can connect to them, and in enterprise scenarios you can daisy-chain switches to cover more connections. How does the switch handle video? What's the maximum resolution and what type of video compression does the switch offer? Do you need sound capability? What about the switch's form factor (is it rack mountable?), the type of cables you'll need for server connections, the maximum number of simultaneous sessions, and the maximum distance the switch allows between the switch and servers? And what kind of failover functionality does it provide? Reliable access to critical resources is a key feature of a KVM over IP platform.

Some switches offer proprietary viewer software for communicating with the switch and others rely on a web browser. If you prefer limited user access to the switch, client software might be best. But get a handle on usability and performance; entry-level products might offer weak security and reliability. If you need to give administrators access regardless of location, use a browser-based interface.

Speaking of security, a major byproduct of the KVM over IP switch's inherent centralization is tighter control of your widespread resources, but the various solutions available today take differing approaches to security. Determine whether the switch takes advantage of your existing authentication technologies or uses its own methods. Does the switch encrypt all signals between itself and managed devices? A great deterrent to intrusion is an encrypted administrative GUI.

Choose Wisely

Web Table 1 shows a listing of the vendors who chose to participate in this year's roundup of KVM over IP switches. You might consider KVM technology basic or elemental, but it's one area where you don't want to choose unwisely.



InstantDoc ID 101689

JASON BOVBERG (jbovberg@windowsitpro.com) is a senior editor for *Windows IT Pro*, *SQL Server Magazine*, and *System iNews*, specializing in networking, hardware, storage/backup, and mobile and wireless. He has 20 years of experience as a writer and editor in magazine, book, and special-interest publishing.

Hit Your IT Bull's-Eye



with **FREE** Trial Software
at Download Central

brought to you by

WindowsIT Pro

Download Central brings you the tools
to meet your most critical IT needs.

A one-stop hub of countless free trial
downloads from leading industry vendors,
Download Central has done all the looking.

All you have to do is see which tool is the
best fit. **And you get to do it all for FREE!**

**Download Central links you
to the solutions you need for:**

- ⊙ Active Directory
- ⊙ Exchange & Outlook
- ⊙ Windows OSs
- ⊙ Desktop Management
- ⊙ SharePoint
- ⊙ SQL Server
- ⊙ Security
- ⊙ Virtualization

Score Your Solution at Download Central!
windowsitpro.com/downloads

WindowsIT Pro

■ Hardware

■ Careers

■ Messaging

INSIGHTS FROM THE INDUSTRY

Intel Pushes Back Tukwila Release, AMD Pushes Phenom II into the Enterprise

Intel is pushing back the release date of its upcoming Itanium processor—currently code-named **Tukwila**—to the middle of 2009 because it's retooling some of the chip's engineering. The additions to Intel's first quad-core Itanium processors include DDR3 memory support and socket-compatibility with future versions of Intel's Itanium chips.

Intel had planned to roll out Tukwila in early 2009, but the design changes are forcing the delay. One of the changes to Tukwila's design is DDR3 memory support. Intel believes users will move to DDR3 sooner rather than later, and the chip maker wants to update its Itanium road map to reflect that belief. The revised Itanium platform will also contain a new piece of memory technology called "scalable buffer memory," which lets OEMs increase the amount of memory the server systems can support. Intel is already supporting DDR3 memory with its processors based on the Nehalem architecture, which came to market in late 2008. Although Intel believes DDR3 is the future, AMD thinks customers want to stick with the less expensive DDR2 for a while. AMD won't switch to DDR3 memory until 2010.

Intel has redesigned Tukwila so that the chip will be socket-compatible with two other Itanium chips that are currently on the road map. Those two Itanium proces-

sors, Poulson and Kittson, are expected to hit the market in the next two to three years. Although Tukwila is built on Intel's 65-nanometer manufacturing process, the company plans to skip 45-nm chips within the Itanium family and move straight to 32-nm chips with Poulson. Intel also plans to roll out mainstream 32-nm processors code-named Westmere in late 2009.

Intel has previously disclosed that Tukwila will offer four processing cores

Although Intel believes DDR3 is the future, AMD thinks customers want to stick with the less expensive DDR2 for a while. AMD won't switch to DDR3 memory until 2010.

and have an initial clock speed of 2GHz. The chip also supports eight instructional threads and uses 30MB of on-die cache. Meanwhile, AMD is rolling out five addi-

tions to its family of **Phenom II** processors. The new Phenom II processors are part of AMD's platform for gaming desktops, code-named **Dragon**, but could have enterprise uses. AMD is putting its high-speed, energy-efficient Phenom II chips up against Intel's Core 2 Duo chips. The chips include a set of energy-efficient tri-core and quad-core chips that AMD seems intent on positioning against the Intel Core 2 Duo E8400 and the Core 2 Quad Q8200.

The rollout comes at a time when AMD has been battling other chip makers, particularly Intel and Nvidia, for market share in the face of declining shipments industrywide. On January 21, AMD confirmed that the prices of some of the new Phenom II processors would be cut by 18 percent.

At the CES event in January, AMD released the Phenom II X4 940 and X4 920 processors as part of touting the speed and energy efficiency of the new Phenom II chips, which range from the X3 710 (2.6GHz) and the X3 720 "Black Edition Processor" (2.8GHz) to the X4 805 (2.5GHz), X4 810 (2.6GHz), and X4 910 (2.5GHz). In addition to supporting newer DDR3 memory, the Phenom II processors will work with DDR2, in a move designed to give AMD's existing partners flexibility.

"In this market, [the enterprise] doesn't get the rapid adoption you see in the consumer side," Dean McCarron, an analyst with Mercury Research, said in an interview. "What I would expect to see happen is corporate clients looking at this technology as the next major refresh opportunity. The next refresh happens right around April, so we'll probably see it show up in June or July."

—Jason Bovberg
InstantDoc ID 101491

Wanted: Your Real-World Experiences with Products

Have you discovered a great product that saves you time and money? Do you use something you wouldn't wish on anyone? Tell the world in a review in What's Hot: Readers Review Hot Products. If we publish your opinion, we'll send you a Best Buy gift card and a free VIP subscription to *Windows IT Pro*! Send information about a product you use and whether it helps you or hinders you to whatshot@windowsitpro.com.



IT Layoffs: What, Me? What Now?

Someone you know is waking up to his or her first day of unemployment. If it's not you, congratulations: You're probably at work sweating bullets trying to keep up with the extra tasks and avoid management's speculating stare. We're with you there. If it is you, we'd like to offer a list of suggestions gleaned from our own experiences.

Don't take it personally. And don't let anyone imply that you could have prevented it. It's tempting to try to analyze why it happened, but even if you come up with a reason, you're not going to change what happened.

Get your fingers walking. Make a list of everything you have to do. Make a list of everything you're terrified will happen. Make a list of everything you have to be grateful for (even if "I'm still breathing" is all you can think of).

Don't withdraw. Get a Facebook account and a LinkedIn account and reach out to professional and personal contacts.

Create a non-industry blurb. For in-person contacts, such as a neighbor at the mailbox, have a blurb ready when they ask how you're doing. Don't lie. Everyone you meet is your new network, so get networking.

Create an industry blurb. For in-person contacts with someone in the industry, have a blurb ready when he or she asks what your experience is in.

Do the math. Obviously you've already thought about your finances. Be as matter-of-fact as you can about your job loss with relatives and friends, even if you think joblessness is a sign of weakness or an admittance of failure. It's not.

Think outside the cubicle. You are not your career. What are the big picture skills you've learned from IT? You know how to assess a situation and determine what the real problem is; you know how to break a huge task into smaller, manageable steps; you know how to visualize something you can't physically see; you have good reasoning and logic skills; you're fairly calm and methodical; you're teachable, adaptable, and can learn new skills and technologies. If the IT jobs were truly drying up, you could do something else—you've got transferrable skills.

Be hopeful. With your sense of humor, your self-confidence, and of course, your experience, you're gold. You'll get through this. We all will.

—Caroline Marwitz
InstantDoc ID 101589

"Change you can Manage"

ChangeVue™ 2009

Download a fully functional 45 day trial from: winitmag.primaltools.com

PrimalScript, ChangeVue, PrimalSQL, PrimalXML, PrimalScope, PrimalMerge, and PrimalPackager are registered trademarks of SAPIEN Technologies, Inc. All other logos, trademarks, and service marks are the property of their respective owners. ©2002-2009 SAPIEN Technologies, Inc. All Rights Reserved.

Troubleshooting Group Policy eLearning series

with
Darren Mar-Elia

WHAT

3 info-packed, online lessons about the inner workings of Group Policy and expert problem-solving techniques.

PLUS live Q&A sessions!

WHEN

June 25, 2009

WHERE

Your computer

COST

\$99

LESSONS

11:00 am EDT – Understanding Group Policy Structure & Processing

12:30 pm EDT – Best Practices for Group Policy Troubleshooting

2:00 pm EDT – Advanced Techniques for Group Policy Troubleshooting

HOW

Register at www.WindowsITPro.com/go/elearning/TroubleShootingGroupPolicy

Step Up Your Group Policy Configuration

Join Group Policy MVP Darren Mar-Elia on June 25, 2009 for break-through lessons on Group Policy.

You'll gain all the skills and tools you need to keep your environment as secured and locked down as you expect it to be.

INSTRUCTOR



Darren Mar-Elia, is President and Chief Technology Officer of SDM Software, Inc. He has more than 20 years combined experience in information technology and software development. Darren has written or contributed to 12 books on Windows management topics, has been a contributing editor at *Windows IT Pro* magazine for more than 10 years.

Learn more about the speaker, sessions,
and how to reserve your seat at:
[www.WindowsITPro.com/go/elearning/
TroubleShootingGroupPolicy](http://www.WindowsITPro.com/go/elearning/TroubleShootingGroupPolicy)

WindowsIT Pro

9 Top IT Skills for 2009

In a Computerworld 2009 annual Forecast survey, IT pros were asked to name the hottest IT skills in 2009. I've noted the key points of each skill in a snapshot format.

1. Programming/application development. SAP, .NET, and C# are cited as the hottest skills in this segment right now. The study estimates that SAP experts make \$35 to \$40 per hour more than average senior technicians.

2. Help desk/technical support. This one shouldn't come as a surprise—with increased outsourcing and American frustration over foreign support staff, having a sharp personality and the ability to explain complex problems simply is in high demand.

3. Project management. Many professionals, despite their experience and savvy, do not have good organization and project management skills. If you are able to acquire these skills and take on a leadership role in projects, you'll be indispensable to your organization.

4. Networking. With the increase in unified communications, there's a high demand for individuals who are knowledgeable in the latest networking technologies.

5. Business intelligence. When it comes to BI, individuals who can understand the systems and collect the right data are

obviously valuable. However, IT pros who can think in terms of business strategy, driving creative ideas for what data to pull and how to use it, are of extreme value.

6. Security. Security threats are abundant and always growing. Organizations not only need someone with a background in security, but also someone who can be proactive and foresee potential threats and eliminate them.

7. Web 2.0. Social networking becomes a bigger part of modern-day business every day, and it's not just limited to Millennials. If you feel like you came to the party too late and won't be able to keep up with the new tools, you're wrong. Most of them are surprisingly intuitive.

8. Data center. Understanding the data center and virtualization is critical, as organizations move to cut energy and storage costs. While many general IT pros are expected to learn these skills, becoming an expert in virtualization will be a smart move.

9. Telecommunications. VoIP, Wi-Fi, WiMAX, Bluetooth—become familiar with these technologies, the devices that are using them, and what growing role they will have in the future.

—Brian Reinholz
InstantDoc ID 101614

Affordable Desktop Management Software

- Software Deployment
- Patch Management
- Asset Management
- Remote Control
- Configurations
- Windows Tools

**20%
Discount**

Starts @ \$545

ManageEngine Desktop Central
www.desktopcentral.com/discount59

- Urgent Desktop Alert

Would you like to alert your users of urgent information direct to their desktops?

Sick of employees not paying attention to important emails?

ACEmessage is a Desktop Alert Solution to deliver messages over Open Applications on the users desktop... instant information!

ACEmessage

- Fast, One-Way, Centralised Distribution
- Multiple Desktop Popup Styles
- Message Thousands of Desktops in Seconds
- Active Directory Integration
- **Free 30 Day Trial**



For more information visit

www.spydaman.com

SMB Communications Pain Points: Is UC the Answer?

A recently released study by SIS International Research, sponsored by Siemens Enterprise Communications, identified the top five pain points in communications for small-to-midsize businesses (SMBs). Part of the objective for the study was to examine SMBs' interest and use of nontraditional technologies and how unified communications (UC) could help companies improve business processes and cut costs.

Top 5 Communications Pain Points

The top pain points identified by this study certainly seem to point to UC as the answer. As reported by the study, here are the top five communications pain points for SMBs:

Inefficient coordination. This points to wasted time setting up meetings and queries sent to the wrong individual because you don't know who has the answer.

Waiting for information. Tied in with inefficient coordination, this latency can delay important business decisions.

Unwanted communications. In addition to spam, think about the time wasted on forwarded joke email messages, or unsolicited sales calls and other low-priority communications.

Customer complaints. This includes time required to deal with negative customer experiences, which can be a result of not being able to reach you in a timely fashion.

Barriers to collaboration. This item includes difficulties establishing collaboration sessions and accessibility problems or not having the right communications tools.

Study respondents reported an average of 17.5 hours a week addressing these pain points, which is equivalent to 40 percent of a 40-hour work week. That's a significant problem for overall productivity—if it's true.

Using such an assumption to get to "a staggering \$524,569" potential loss annually for SMBs seems like blatant scare tactics to scare up business for Siemens' UC solutions.

Problems with the Study

Here's where my cynical side kicks in. As Homer Simpson says, "Facts are meaningless. You can use facts to prove anything that's even remotely true." Or, to put it another way, how much can you trust a sponsored report?

I often wonder how many of these sponsored studies never get released because the results don't match the marketing message of the sponsoring organization. And even with a study such as this one that does get released, you have to watch out for the organizational spin.

For example, in the Conclusions section of the study, it's reported that the cost of these pain points could be "\$5,246 per year per employee, assuming that 100% of the time reported addressing these issues is unproductive. Thus, for example, a SMB with 100 employees could be leaking a staggering \$524,569 annually as a result of inefficiencies in communication."

The problem here is with that assumption of complete unproductivity. Yes, you might spend 3.5 hours a week waiting for information, but it seems unreasonable to say that you're not doing something else productive during that time. Using such an assumption to get to "a staggering \$524,569" potential loss annually for SMBs seems like blatant scare tactics to scare up business for Siemens' UC solutions.

UC Could Be the Answer

Don't get me wrong—I've read a lot about UC that really does impress me about its usefulness and potential cost savings to businesses. (And I can still remember the joy I felt the first time I got a voicemail message through web mail while working from home.) But I'm suspicious of "studies" that sound like marketing. Download the full study at tinyurl.com/cst5bd to judge for yourself.

—B. K. Winstead
InstantDoc ID 101585

Prime Your Mind

with Resources from Left-Brain.com

Left-Brain.com is the newly-launched online superstore stocked with resources focused on meeting the needs of IT professionals like you.

Featured Product:
PowerShell 101: A Quick-Start Guide to PowerShell
by Robert Sheldon

Ease your scripting pains with the flexibility of PowerShell. Learn how to use PowerShell to perform various tasks with this guide's 6 introductory lessons—complete with helpful figures, expert explanations and detailed code.

Order your downloadable eBook today for only \$15.95!

*Plus shipping and applicable tax.

www.left-brain.com

Get Full Access That Fits Your Schedule

with the
WindowsITPro

Monthly Online Pass


For only \$5.95 per month, you'll become a member of the world's largest independent IT community and get access to over 10,000 Windows IT Pro online articles! Plus, you'll get the latest digital issue of Windows IT Pro magazine!

Sign up today to start getting the answers you need when you need them.

www.windowsitpro.com/go/MonthlyPass

Are Your IIS Servers Under Attack?

Block all unwanted IIS traffic with ThreatSentry



download free trial

- IIS web application firewall & IPS
- stops known, new and internal threats
- blocks sql injection, xss, dos and more
- reinforces regulatory compliance



ITV/Software Solutions
Data Management Solutions

sales@privacyware.com • www.privacyware.com • 732.212.8110 x235

Making a Server Move?

Don't sweat the small stuff...or the big stuff.

MigratePro handles the grunt work for you; migrating your shares, share settings, and data to your new server.



Download a Free 30-day Trial

Details at www.SoftwarePursuits.com

800-367-4823 or 650-372-0900 **SoftwarePursuits**

Prime Your Mind

with Resources from Left-Brain.com



Left-Brain.com is the newly-launched online superstore stocked with resources focused on meeting the needs of IT professionals like you.

Featured Product:

PowerShell 101: A Quick-Start Guide to PowerShell

by Robert Sheldon

Ease your scripting pains with the flexibility of PowerShell. Learn how to use PowerShell to perform various tasks with this guide's 6 introductory lessons—complete with helpful figures, expert explanations and detailed code.

Order your downloadable eBook today for only \$15.95*!



*Plus shipping and applicable tax.

www.left-brain.com

WindowsITPro

For detailed information about products in this issue of *Windows IT Pro*, visit the Web sites listed below.

COMPANY/URL	PAGE	COMPANY/URL	PAGE
AdventNet ManageEngine	60	Software Pursuits Inc.	62
www.desktopcentral.com/discount59		www.SoftwarePursuits.com/MPro25	
AvePoint Inc.	14, 44, 49	Special Operations	14
www.avepoint.com		www.specopssoft.com/gold	
Bid2Learn	25	Spydaman Design Studios Ltd.	60
www.Bid2Learn.com		www.spydaman.com	
Diskeeper Corporation	8B	St. Bernard Software	Cover 4
www.diskeeper.com		www.stbernard.com	
Fpweb.net	49	Sunbelt Software Inc.	7
www.fpweb.net		www.TestDriveVipre.com	
GFI Software Ltd.	Cover tip	Train Signal	Cover 3
www.gfi.com/winitpro		www.trainsignal.com	
Microsoft Corporation	21	Unisys	16B, 32B
www.SQLServerEnergy.com		www.unisys.com	
Microsoft Corporation	Cover 2	Windows Connections 2009	52
www.microsoft.com/virtualization		www.WinConnections.com	
Netikus	18	Windows IT Pro	42, 56, 59, 62
www.eventsentry.com		www.windowsitpro.com	
Privacyware	62		
www.privacyware.com			
R1soft	26		
www.r1soft.com			
Sapien Technologies	58		
www.sapien.com			

VENDOR DIRECTORY

The following vendors or their products are mentioned in this issue of *Windows IT Pro* on the pages listed below.

AMD	57
Apple	51
Citrix	3, 51
Cortex I.T.	51
HP	3
Intel	57
Mitsubishi International Corporation	63
Remote Access Software	51
Siemens Enterprise Communications	61
Unbounded Solutions	64
VMware	3
Wyse Technology	3
Xobni	53

DIRECTORY OF SERVICES | WINDOWS IT PRO NETWORK

Search our network of sites dedicated to hands-on technical information for IT professionals.

www.windowsitpro.com

Support

Join our discussion forums. Post your questions and get advice from authors, vendors, and other IT professionals.

www.windowsitpro.com/forums

News

Check out the current news and information about Microsoft Windows technologies.

www.wininformant.com

EMAIL NEWSLETTERS

Get free NT/2000/XP/2003 news, commentary, and tips delivered automatically to your desktop.

[Exchange & Outlook UPDATE](#)

[Scripting Central](#)

[Security UPDATE](#)

[SQL Server Magazine UPDATE](#)

[ToTheSharePoint Newsletter](#)

[WindowsDevPro UPDATE](#)

[Windows IT Pro UPDATE](#)

[Windows Tips & Tricks UPDATE](#)

[WinInfo Daily UPDATE](#)

www.windowsitpro.com/email

RELATED PRODUCTS

Custom Reprint Services

Order reprints of *Windows IT Pro* articles. Contact Joel Kirk at jkirk@penton.com.

Super CD/VIP

Get exclusive access to all of our print publications, including *Windows IT Pro*, via the new, banner-free VIP Web site.

www.windowsitpro.com/sub/vip

Article Archive CD

Access every article ever printed in *Windows IT Pro* magazine since September 1995 with this portable and speedy tool.

www.windowsitpro.com/sub/cd

SQL SERVER MAGAZINE

Explore the hottest new features of SQL Server, and discover practical tips and tools.

www.sqlmag.com

ASSOCIATED WEB SITES

WindowsDev Pro

Discover up-to-the-minute expert insights, information on development for IT optimization, and solutions-focused articles at WindowsDevPro.com, where IT pros creatively and proactively drive business value through technology.

www.windowsdevpro.com

Office & SharePoint Pro

Dive into Microsoft Office and SharePoint content offered in specialized articles, member forums, expert tips, and Web seminars mentored by a community of peers and professionals.

www.officesharepointpro.com

NEW WAYS TO REACH

WINDOWS IT PRO EDITORS:

LinkedIn: To check out the *Windows IT Pro* group on LinkedIn, sign in on the LinkedIn homepage (www.linkedin.com), select the Search Groups option from the pull-down menu, and use "Windows IT Pro" as your search term.

Facebook: We've created a page on Facebook for *Windows IT Pro*, which you can access at: <http://tinyurl.com/d5bquf>. Visit our Facebook page to read the latest reader comments, see links to our latest web content, browse our classic cover gallery, and participate in our Facebook discussion board.

Twitter: Visit the *Windows IT Pro* Twitter page at www.twitter.com/windowsitpro.

Regional Forums: We've introduced regional areas in our online forums, allowing IT user group leaders and other readers interested in meeting locally to more easily communicate with each other. Visit our forums at www.windowsitpro.com/forums and scroll down to see the new regional forums.

Windows IT Pro

TOP 10 Unintentionally Hilarious Domain Names

These unfortunate URLs have been floating around the Internet for a long time, but it's worth noting that as of this writing, they're all still happily active and hilarious.

10. TeachersTalk (www.teachertalk.co.uk)
9. DollarsExchange (www.dollarsexchange.com)
8. TIE: AccessTherapist (www.accesstherapist.com) and Therapist Finder (www.therapistfinder.com)
7. WhoRepresents? (www.whorepresents.com)
6. Italian Power Generator company (www.powergenitalia.com)
5. GoTahoe (www.gotahoe.com)
4. World Taekwondo Federation (www.wtf.org)
3. TIE: Pots of Art (www.potsofart.com) and Speed of Art (www.speedofart.com)
2. IPAnywhere (www.ipanywhere.com)
1. Pen Island (www.penisland.net)



It's never human error, of course

Not an early riser

User Story of the Month

On weekend duty, working for the County Library District, I received a service call from a major branch. Patrons couldn't log on to public computers. This branch was the only one experiencing a problem, so I thought it might be the onsite schedule server. I called and asked a staffer to reboot the machine. He performed the reboot, but the system halted repeatedly during the POST sequence. I left immediately, driving the 30 miles to the branch to examine the machine. When I arrived, I reached for the keyboard, which someone has leaned against the wall; I was about to hit Ctrl+Alt+Del and watch the POST. However, as soon as I moved the keyboard, POST resumed and the machine finished booting. Apparently, a technician had needed some room on the table and had set the keyboard against the wall so that the Pause key was pressed. This situation had no repercussions until the nightly reboot.

—Johnny Reel



May 2009 issue no. 177, *Windows IT Pro* (ISSN 1552-3136) is published monthly. Copyright 2009, Penton Media, Inc., all rights reserved. Windows is a trademark or registered trademark of Microsoft Corporation in the United States and/or other countries, and *Windows IT Pro* is used under license from owner. *Windows IT Pro* is an independent publication not affiliated with Microsoft Corporation. Microsoft Corporation is not responsible in any way for the editorial policy or other contents of the publication. *Windows IT Pro*, 221 E. 29th St., Loveland, CO 80538, (800) 793-5697 or (970) 663-4700. Sales and Marketing Offices: 221 E. 29th St., Loveland, CO 80538. Advertising rates furnished upon request. Periodicals Class postage paid at Loveland, Colorado, and additional mailing offices. POSTMASTER: Send address changes to *Windows IT Pro*, 221 E. 29th St., Loveland, CO 80538. SUBSCRIBERS: Send all inquiries, payments, and address changes to *Windows IT Pro*, Circulation Department, 221 E. 29th St., Loveland, CO 80538. Printed in the USA. BPA Worldwide Member.

**NEED TO SECURE YOUR JOB, WANT TO START A NEW CAREER...
WITHOUT BREAKING YOUR PIGGY BANK?**



YOU CAN WITH TRAIN SIGNAL TRAINING COURSES FOR UNDER \$400!

Our Computer Training Software:

- Is Scenario Based! It Mirrors Real World Challenges and Lays a Solid Foundation for Your Career!
- Helps You Prepare for Your Certification Exam the Right Way!
- 90-Day Money Back Guarantee!

We offer training in:



Get a headstart and call our toll free number today!
(888) 229-5055



TRAINSIGNAL

THE GLOBAL LEADER IN PROFESSIONAL COMPUTER TRAINING™

www.trainsignal.com

Copyright © 2002-2009 Train Signal, Inc. All Rights Reserved. All logos and trademarks are property of their respective owners.

From: Renewal time, here comes
the pain again

To: Predictable pricing &
consistent support

NO-NONSENSE WEB FILTERING

That's what you'll get when you switch to iPrism from St Bernard – the award-winning web filter that's easier in every way, and less expensive to own.

iPrism is changing the way companies and schools everywhere handle their web filtering. With blazing throughput speeds up to 100+ Mbps, anti-virus protection and seamless XenApp and Active Directory integration, iPrism is the appliance-based solution of choice for customers and institutions of any size.

Find out more about the easiest-to-deploy, most highly rated web filtering solution ever – the industry's ONLY Citrix-ready web filtering appliance.

Call 1.800.782.3762 or go to www.SwitchToiPrism.com/flip



FLIP THE SWITCH

Get your **FREE** iPrism® Switch Kit today:

FREE 30-day onsite evaluation

that can be deployed without any client or network changes

FREE enhanced technical support

for setting up matching policies, reports & alerts based on your current settings

INCENTIVE PRICING & A FREE T-SHIRT

just for watching a live demo



iPrism® h-Series, the world's #1 Web Filtering appliance.

© 2009 St Bernard Software, Inc.